Hardware Overview

Hardware is the physical part of computing. From disk drives to printer cables, from speakers to printers, hardware is the part of computing you can pick up, move around, open, and close. Although hardware might represent the glamorous side of computing (whose computer is faster, has a larger hard disk, more memory, and so on), it can do nothing without software and firmware to provide instructions. Hardware failures can take place because of loose connections, electrical or physical damage, or incompatible devices.

Software Overview

Software provides the instructions that tell hardware what to do. The same computer system can be used for word processing, gaming, accounting, or Web surfing by installing and using new software. Software comes in various types, including operating systems, application programs, and utility programs.

Operating systems provide standard methods for saving, retrieving, changing, printing, and transmitting information. The most common operating systems today are various versions of Microsoft Windows. The 2003 version of the A+ Certification Exam focuses on all recent 32-bit desktop versions of Windows (Windows 9x and Me, Windows NT 4.0 Workstation, Windows 2000 Professional, and Windows XP). Because operating systems provide the "glue" that connects hardware devices and applications, they are written to work on specified combinations of CPUs and hardware.

Operating system commands come in two major types: internal and external. Internal commands are those built into the operating system when it starts the computer. External commands require that you run a particular program that is included with the operating system.

Application programs are used to create, store, modify and view information you create, also called **data**. Because an operating system provides standard methods for using storage, printing, and network devices to work with information, applications must be written to comply with the requirements of an operating system and its associated CPUs.

A+ Certification does not require any knowledge of application programs, but to provide the best technical support, you should learn the basics of the major applications your company supports, such as Microsoft Office, Corel WordPerfect Office, Adobe Photoshop, and many others.

Certifications are available for major operating systems and applications, and seeking certifications in these areas can further improve your chances of being hired and promoted.

Utility programs are used to keep a computer in good working condition or to set up new devices. In the operating system chapters, you'll learn how to use the major utilities that are included with Windows.

Because these utilities have limited capabilities, you might also want to invest in other utility programs, such as Symantec's Norton System Works or PowerQuest's Drive Image or Partition Magic, for use in your day-to-day work; however, only standard Windows utilities, such as ScanDisk, Fdisk, Disk Management, Defrag, and others, are covered on the A+ Certification Exam.

Firmware Overview

Firmware represents a middle ground between hardware and software. Like hardware, firmware is physical: a chip or chips attached to devices such as motherboards, video

cards, network cards, modems, and printers. However, firmware is also software: Firmware chips (such as the motherboard BIOS) contain instructions for hardware testing, hardware configuration, and input/output routines. In essence, firmware is "software on a chip," and the software's job is to control the device to which the chip is connected. Because firmware works with both hardware and software, changes in either one can cause firmware to become outdated. Outdated firmware can lead to device or system failure or even data loss.

Until the mid-1990s, the only way to change firmware was to remove the chip and replace it with one containing new instructions. Most firmware today is flashable, meaning that its contents can be changed through software. You'll learn more about the most common type of firmware, the motherboard BIOS, in Chapter 6, "BIOS and CMOS Configuration."

Memory: RAM and ROM

There are two types of **memory** in a computer: RAM and ROM. In the early days of personal computing, some vendors made their systems sound more impressive by adding these two totals. However, their function in computer systems is verydifferent.

The contents of RAM (random access memory) can be accessed in any order and can change instantly. They are in constant flux as you start a computer; load its operating system and drivers for particular devices; load an application; create, store, change, and copy data; and shut down the computer.

Programs are loaded into RAM; until data is stored, it exists only in RAM (that's why you should save your work so often!). The "enemies" of data stored in RAM include

- System crashes and lockups
- User error (forgetting to save before you close a program)
- Power failures

Because most types of RAM must receive a steady dose of electricity to keep its contents around, even momentary power failures can destroy its contents. Because all data must be created or changed in RAM before it's stored, you must make sure that RAM is working correctly. In Chapter 7, "RAM," you'll learn more about adding, configuring, and using RAM.

Flash memory, also called Flash RAM, is a special type of RAM that uses electricity to change its contents, but doesn't require electricity to maintain its contents. It is used in BIOS chips and in storage for digital music players and digital cameras.ROM stands for read-only memory, meaning that although its contents, like those of RAM, can be accessed in any order, ROM's contents can't be changed by normal computer operations. So, what is ROM good for?

Because ROM's contents don't change when a system is powered down or restarted, it's the perfect storage place for firmware. As we saw earlier, firmware is the "software on a chip" used to control various devices in the computer. ROM isn't suitable for software storage, however, because its capacity is too limited for today's large programs. And, of course, ROM can't be used to store data files that are constantly changing.

The way that ROM chips have been made has changed several times over the years. Originally, ROMs contained a permanently etched pattern; later, ROMs were made of reprogrammable materials that could be changed through controlled ultraviolet light or electricity. Because the chip had to be removed from the motherboard for replacement or reprogramming, changing the contents of ROMs was difficult and inconvenient. If you

think that opening up a single system to change its ROM chip is a pain, imagine performing the same job on dozens of PCs!

Fortunately, current ROMs can be reprogrammed with software. This process is called **flashing** the ROM and is performed with the BIOS firmware found on motherboards and in modems, among other devices. Want to learn more? The reasons for upgrading ROMs and the methods used are covered in Chapter 6.

Dissecting Your Computer

Before you can troubleshoot your computer, you need to understand what a computer is and how it's put together. A computer (or PC, for personal computer) is not a single unit, but is instead a collection of hardware subsystems including

- Video
- Storage
- Input devices
- Printers and other output devices
- Audio
- Networking
- Processor
- Memory
- Power
- Motherboard

Whenever a computer stops working, you can trace it back to the failure of one or more of these subsystems (the device and its cables). What controls the subsystems?

These subsystems are controlled by two types of software:

- A system BIOS (basic input/output system) chip on the motherboard (an example of "software on a chip," or firmware)
- The operating system and its device drivers (files that tell Windows how to use your PC's hardware)

In the following sections, I use the term *point of failure* to refer to a component or BIOS configuration that could cause problems for your system. This term isn't meant to suggest that computers are constantly on the verge of having a problem, but that some parts of the computer are more likely to cause problems than others.Hardware and operating system software are both used by application programs such as Microsoft Office, Adobe Photoshop, Quicken, and many, many others to create, change, store, print, and transmit information.

With even the simplest devices and software depending upon so many other factors, troubleshooting your computer can be a challenge. But, if you don't know the details of what's inside a typical computer and how all this hardware and software relates to each other, it's just about impossible.

This chapter introduces you to the major components you will find in typical computers, including those prone to being a point of failure. Think of it as an anatomy lesson, but without the formaldehyde or nasty smells.

The Outside Story of Typical Computers

Although the "inside" story of computers is even more complicated than the outside, don't neglect taking a good look at the outside of your system when it's time to troubleshoot a

computer problem. You also need to check out the outside of the computer when it's time to plug in a peripheral to see if there's a suitable connector for it.

The outside of the computer is where you'll find

- Cable connections for external peripherals such as cable modems, printers, monitors, and scanners
- Drive bays for removable-media and optical drives
- The power supply fan and voltage switch
- The power cord between the computer and the wall outlet (and between external peripherals such as monitors and printers and the wall outlet)
- The power switch, reset button, and signal lights

All in all, the outside of the computer is a good place to start when your computer has stopped working, even if you're not sure if that's where the problem lies.

The Front View of a Typical Desktop Computer

A typical "desktop" computer actually sits on the floor in most offices, and resembles Figure 2.1.

Figure 2.1. The front panel of a typical desktop computer with some higher-end internal peripherals. The components seen here vary from PC to PC.

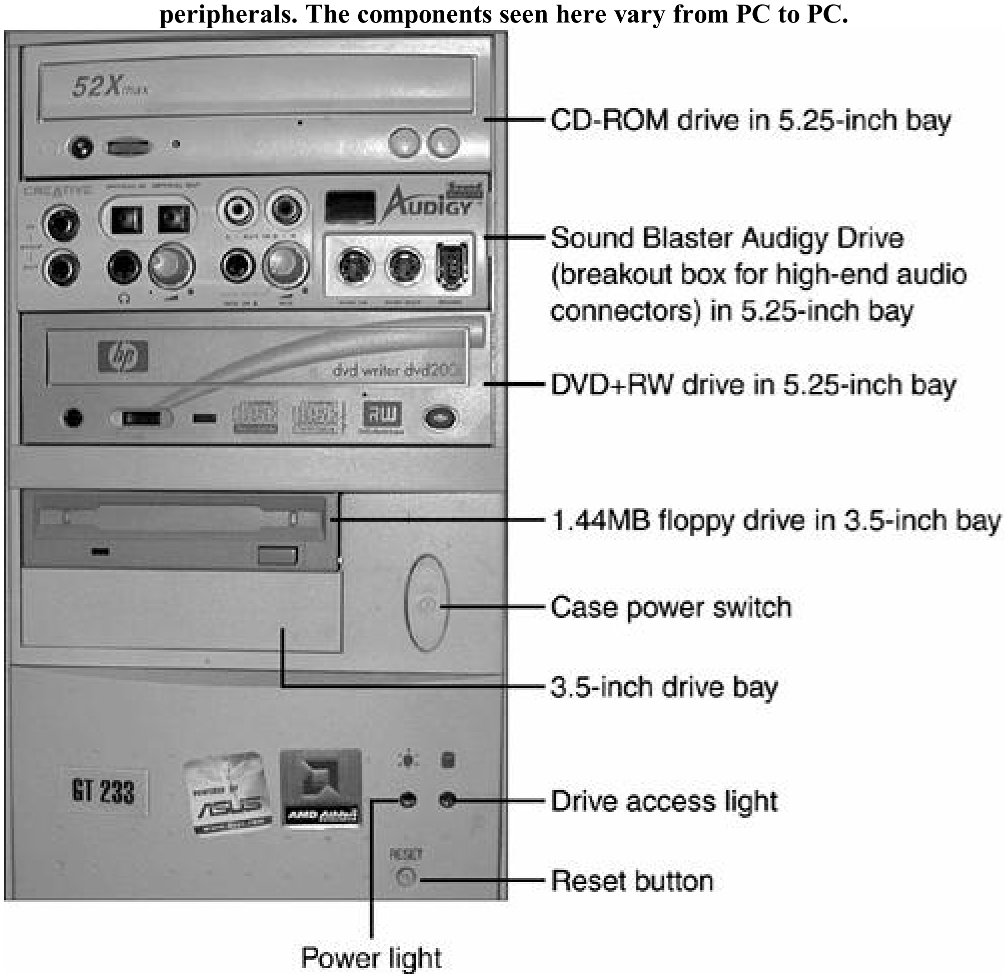

The case shown in Figure 2.1 is sometimes referred to as a mid-tower case. This computer has room for up to six internal drives (three in 5.25-inch bays and three in 3.5-inch bays; one of the 3.5-inch bays can't be seen in Figure 2.1).

CD-ROMs, CD-RWs, DVD-ROMs, and similar optical drives, as well as large removable-media drives, such as the Castlewood Orb and Iomega Jaz, use 5.25-inch drive bays. Floppy drives, hard drives, and smaller removable-media drives, such as Iomega Zip and SuperDisk LS-120/LS-240, use 3.5-inch drive bays. 3.5-inch drives also can be installed in 5.25-inch drive bays with special mounting brackets.

As Figure 2.1 also shows, you can add other types of devices to the 5.25-inch drive bays, such as the Sound Blaster Audigy Drive (a breakout box for this popular sound card), cooling fans, drive-selection switches, and front-mounted connectors for IEEE-1394 and USB ports.

Note that the Sound Blaster Audigy Drive has connectors for some types of speakers and other multimedia devices. Some cases also feature built-in USB, serial, and IEEE-1394a ports. Unfortunately, front-mounted ports are still quite rare. Most devices still connect to the rear of the computer, which makes it more difficult to fix problems caused by loose cables and gets your knees dusty.

Points of Failure on the Front of a Computer

The front of the computer might provide valuable clues if you're having problems with a system. In case of problems, check the following common points of failure for help:

- Can't read CD or DVD media— The drive door on the CD-ROM or other optical drive might not be completely closed or the media might be inserted upside down; press the eject button to open the drive, remove any obstacles, reseat the media, and close the drive.

  tip
- You also can eject optical media with Windows Explorer under My Computer. Right-click on the drive and select Eject. If the drive doesn't eject the media, there could be a problem with the drive's data cable, cable connection, or power connection.Can't shut down the computer with the case power switch— The case power switch is connected to the motherboard on ATX and Micro-ATX systems, not directly to the power supply as with older designs. The wire might be loose or connected to the wrong pins on the motherboard. Keep in mind that most systems require you to hold in the power button for about four seconds before the system will shut down. If the computer crashes, you might need to shut down the computer by unplugging it or by turning off the surge suppressor used by the computer. Some ATX power supplies have their own on-off switches, but this is not commonplace.
- Can't see the drive access or power lights— As with the case power switch, these lights are also connected to the motherboard. These wires might also be loose or connected to the wrong pins on the motherboard.
- Can't use USB, IEEE-1394, or digital camera (serial) ports on the front of the system— Some systems have these ports on the front of the computer as well as the rear. Front-mounted ports are connected with extension cables to the

motherboard. If the cables inside the case are loose, the ports won't work. If the ports are disabled in the system BIOS, the ports won't work.

The term desktop computer is a bit misleading today because few computers actually sit on the user's desk anymore (when I started working with computers in 1983, real desktop computers were virtually all there were). However, this term survives to describe computers that Use standard internal components such as motherboards, processors, memory, drives, sound cards, and video cards Can be upgraded and rebuilt by the user without special tools Use separate input devices (keyboard, mouse, or other pointing device) Most (but not all) desktop computers are multipiece units with a separate keyboard and monitor, although the iMac has inspired a few compact "all-in-one" PCs, which incorporate a monitor.As you can see from this section, in many situations, you will need to open the case to resolve a problem, even though the symptoms might first manifest themselves outside the computer. To learn more about the typical components found in a PC, see "Inside a Typical PC ," later in this chapter.

The Rear View of a Typical Desktop Computer

If the video has gone missing in action, you can't connect to the Internet, or your printer refuses to print, it's time to check out the rear of the computer.

Figure 2.2 shows the rear of a typical desktop computer when common peripheral cables are attached, while Figure 2.3 shows the rear of the same computer after the cables are removed.

Figure 2.2. The rear panel of a typical desktop computer with common external peripheral cables attached.

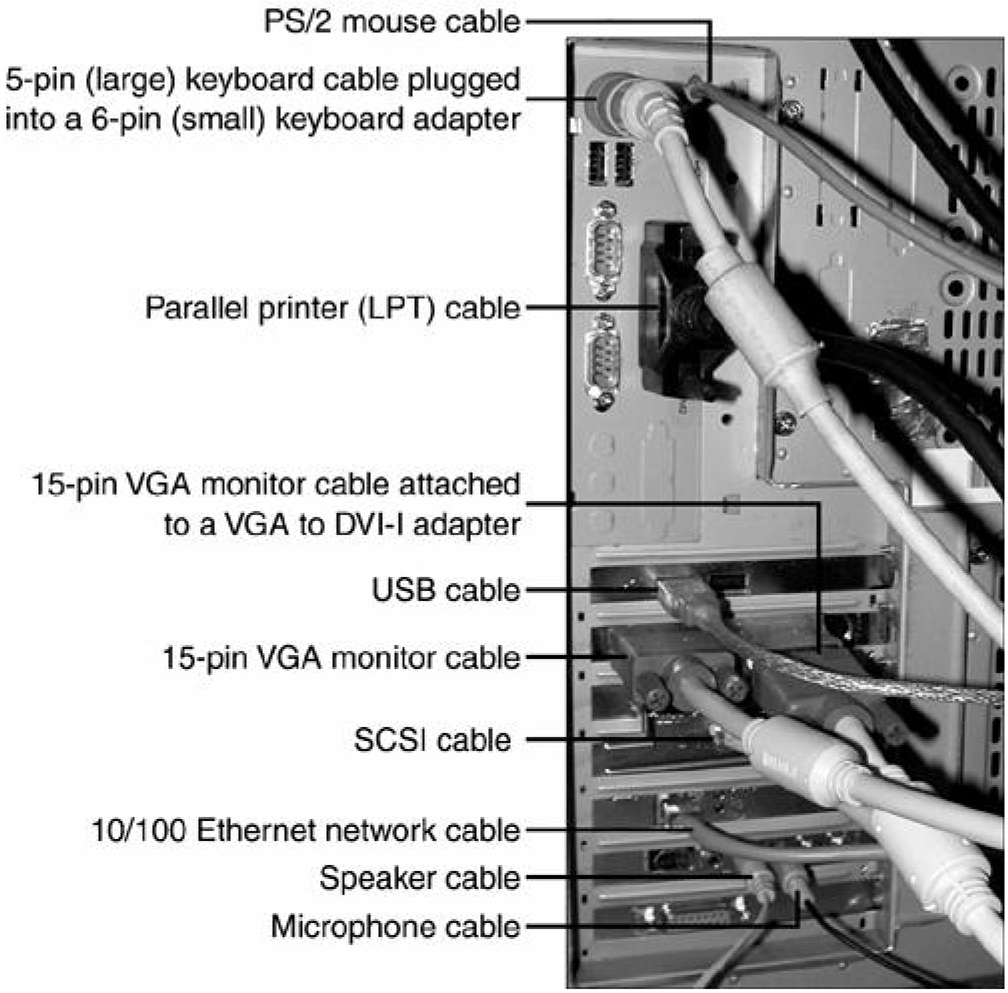

Figure 2.3. The rear panel of a typical desktop computer with built-in ports (top) and ports on add-in cards (bottom).

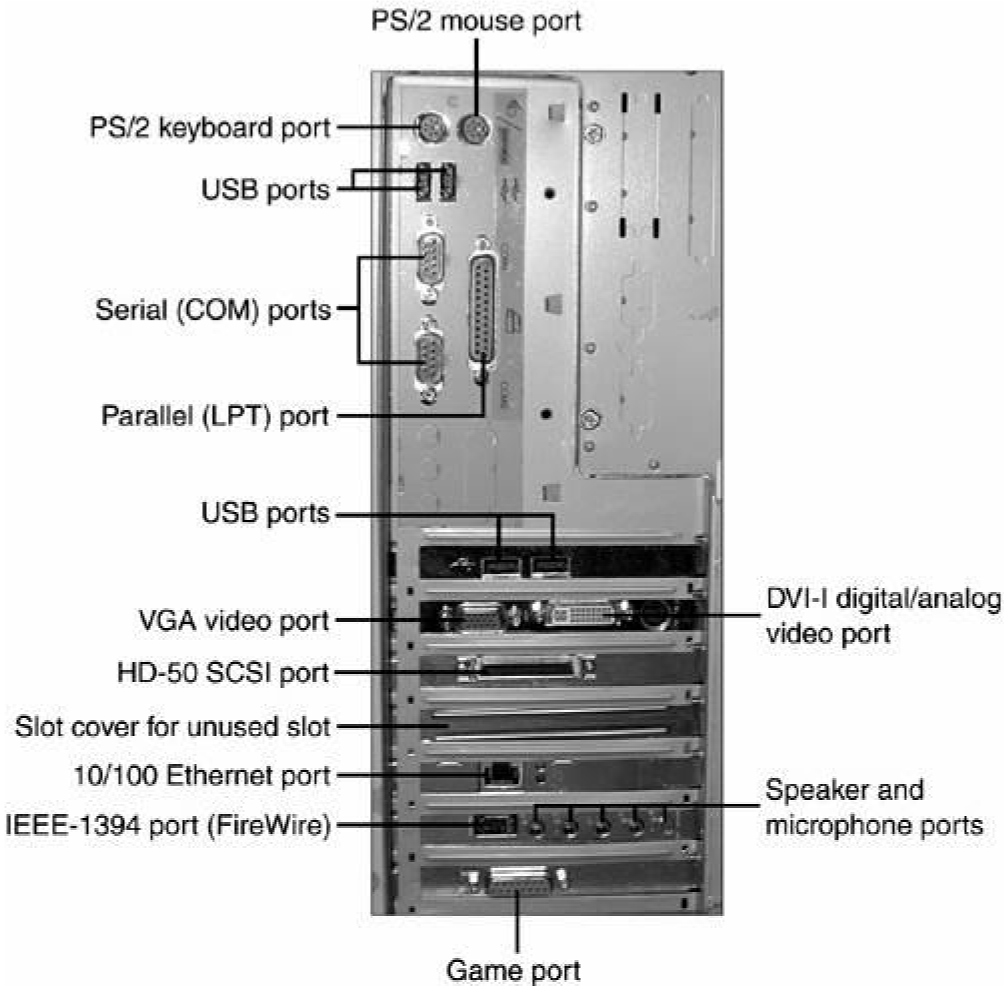

Figures 2.2 and 2.3 are typical of computers built in local computer shops (also called "box shops") from standard parts. However, the rear panel of a computer purchased in a retail store or sold as a corporate workstation might more closely resemble the rear panel shown in Figure 2.4. These computers often have built-in sound and VGA video ports (visible in Figure 2.4), and some might also include built-in 10/100 Ethernet ports.

Figure 2.4. The rear panel of a desktop computer with built-in VGA video and sound, and the standard peripheral cables.

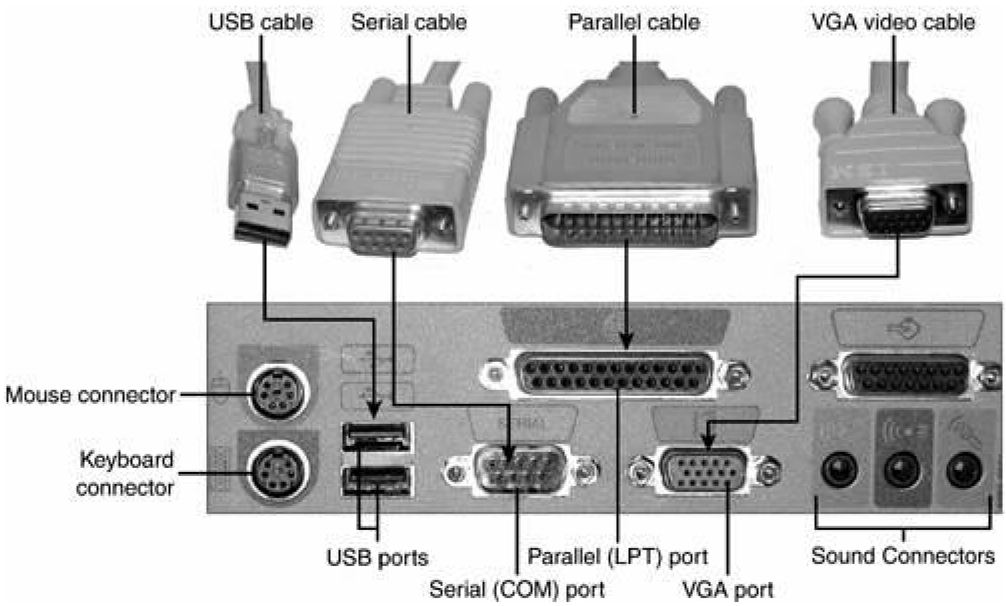

The computers pictured in Figures 2.2–2.4 use motherboards that are members of the ATX family of motherboard designs. The ATX design, and its more compact siblings, Micro-ATX and Flex-ATX, are found in virtually all computers sold since 1997; the ATX family is far and away the most common industry–standard motherboard form factor used today.

tip

Although the cables shown in Figure 2.4 are distinctly different from one another, some cables (such as speaker and microphone cables) use the same type of connector. How can you tell what's what? Most recent computers, motherboards, and peripherals use the PC99 color coding standards. See them online at <http://www.pcdesguide.org/documents/pc99icons.htm>.What are the differences between ATX, Micro-ATX, and Flex-ATX? ATX motherboards have six or seven expansion slots inside; Micro-ATX motherboards, which are used most frequently in low-cost retail-store computers sold for home use, normally have three or four expansion slots. Flex-ATX motherboards have no more than two expansion slots, and often have no expansion slots at all; they are used most often in small form-factor and slimline PCs used as corporate workstations.

Older types of computers use different types of motherboards that have built-in ports mounted under the expansion slots (LPX) or have built-in ports wired to connectors mounted in the expansion slots (Baby-AT). See Chapter 4 , "The Motherboard and CPU," for more information about the ATX family, Baby-AT, and LPX motherboard designs.

Above the rear panel details shown in Figures 2.2–2.4 is the power supply (see Figure 2.5). Power supplies actually convert high-voltage AC power into the low-voltage DC power used inside the computer. Because conversions of this type create heat, the power supply has a fan to cool itself and also help overall system cooling.

Figure 2.5. A typical power supply mounted in a computer.

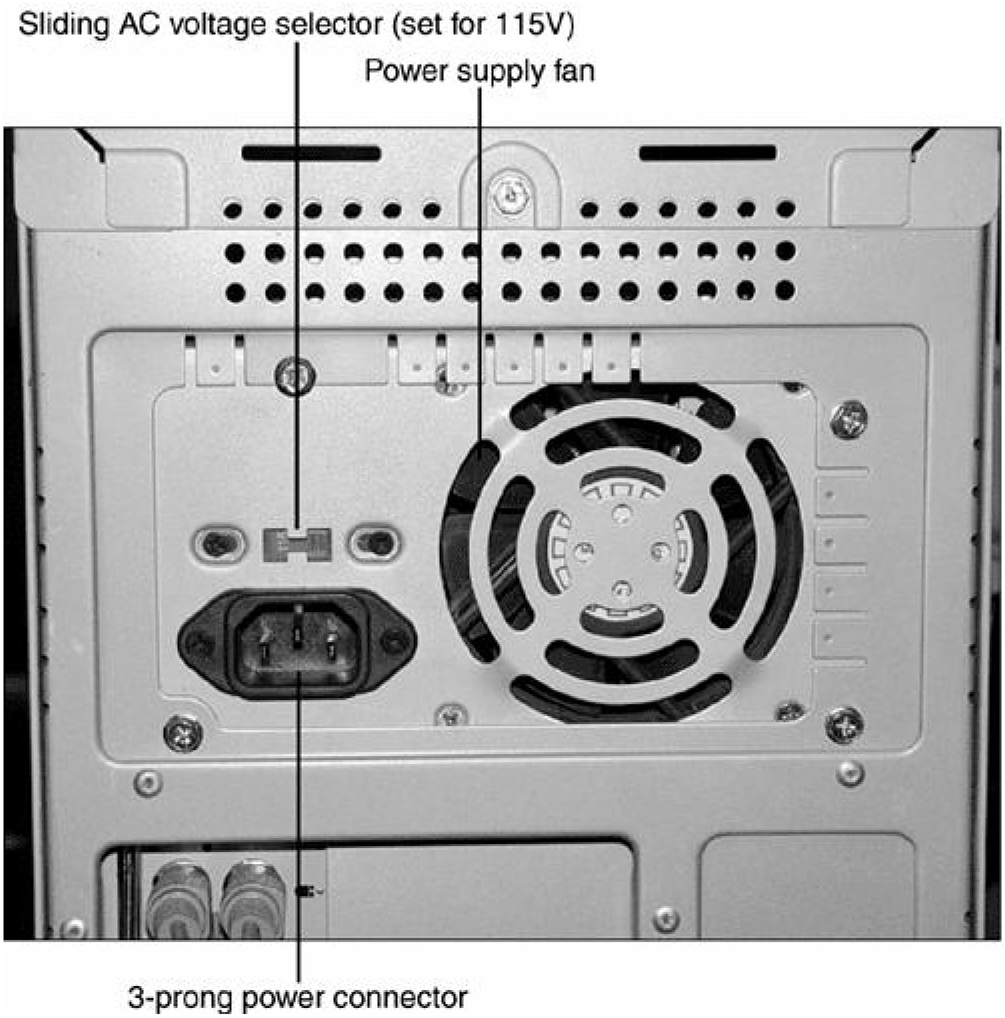

Although a few power supplies can switch between 115-volt and 230-volt (V) services automatically, most use a sliding switch.

Learn more about the power supply in Chapter 5 , "Power Supplies and Circuit Testing."

The Parallel Port

The parallel port, previously shown in Figure 2.4, is one of the oldest multipurpose interfaces found on a typical PC. The parallel port uses a DB-25F connector, and was originally designed for use with parallel printers. However, it has since been adapted for use with removable-storage drives of varying types (Zip, SuperDisk, CD-ROM, CD-RW), scanners, and for direct parallel (Direct Cable Connection or LapLink) file transfer. The parallel port is included as part of most ATX and Micro-ATX computers' rear panels as shown in Figure 2.4, but it can also be built into an add-on card, which can be installed into the PCI or ISA expansion slots on the motherboard.

Don't be surprised if you work on PCs without parallel ports. The USB port can do every job the parallel port can do, and so-called "legacy-free" computers don't include built-in parallel ports.

If the parallel port is built into the motherboard, as in Figure 2.4, it is considered an ISA device and its IRQ (normally IRQ 7) cannot be shared with other devices. For more details

on the hardware resources used by the parallel port, see "Hardware Resources ," later in this chapter.

The Serial Port

The serial port, shown in Figure 2.4, was introduced along with the parallel port on the first PCs. The original serial ports used a DB-25M connector, but virtually all serial ports in recent years have used a DB-9M connector as seen in Figure 2.4. The serial port was originally intended for use with analog (dial-up) modems and serial printers, but has been most often used in more recent times for data transfer with programs such as Direct Cable Connection or LapLink, battery backup (UPS/SPS) signaling, PDA docking station connections, or connections for pointing devices such as mice.

Like the parallel port, the serial port is also being replaced by the USB port. Although most recent computers still have one or two serial ports built into the motherboard, legacy-free computers don't have any serial ports. Serial ports can also be added to the ISA or PCI expansion slots on the motherboard.

If the serial port is built into the motherboard, it is an ISA device and it cannot share its IRQ with devices other than serial ports. Although COM 1 and COM 3 "share" IRQ 4, and COM 2 and COM 4 "share" IRQ 3, in practice this means only that COM 1 or COM 2 can be used only when COM 3 or COM 4 is not being used, and vice versa. For more details on the hardware resources used by the serial port, see "Hardware Resources ," later in this chapter.

The SCSI Interface

Here's an interface you won't see nearly as often as serial or parallel ports: SCSI. SCSI cards are installed only when a particular device needs to connect via SCSI.

The SCSI interface can have several forms, depending upon the type of SCSI card installed in your computer (it's rare to have a motherboard with built-in SCSI). The most common types of SCSI connectors are shown in Figure 2.6.

Figure 2.6. Typical SCSI interfaces: HD-68 (top), HD-50 (second from top), DB25 (third from top), and Centronics 50 (bottom).

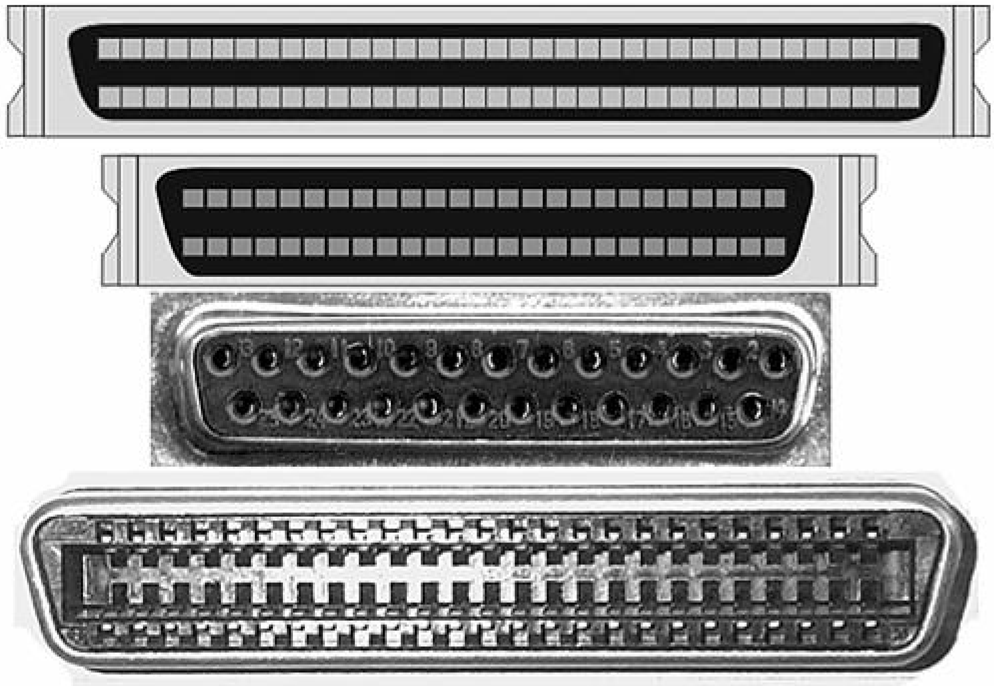

The HD-68 pin connector is used for high-performance Wide SCSI devices such as hard drives and tape backups. The other three interfaces are used for various types of Narrow SCSI devices, including CD-ROM, CD-R, and CD-RW drives; scanners; and removable-media drives. Most SCSI cards can support internal and external devices.

The 25-pin connector shown in Figure 2.6 is found primarily on low-cost SCSI cards bundled with scanners or Zip drives. It uses the same DB-25F connector used by the parallel port, but is not interchangeable. To determine if a 25-pin adapter on the rear of your system is a SCSI port, use these methods:

- A 25-pin SCSI connector will usually say SCSI on the card bracket at the rear of the computer or be marked with the SCSI SE SCSI symbol shown in Figure 2.7.

Figure 2.7. Logos used to mark SCSI ports and cables. SE signaling is used by narrow SCSI devices. LVD signaling is used by most Wide SCSI devices. LVD/SE ports can work with LVD or SE devices.

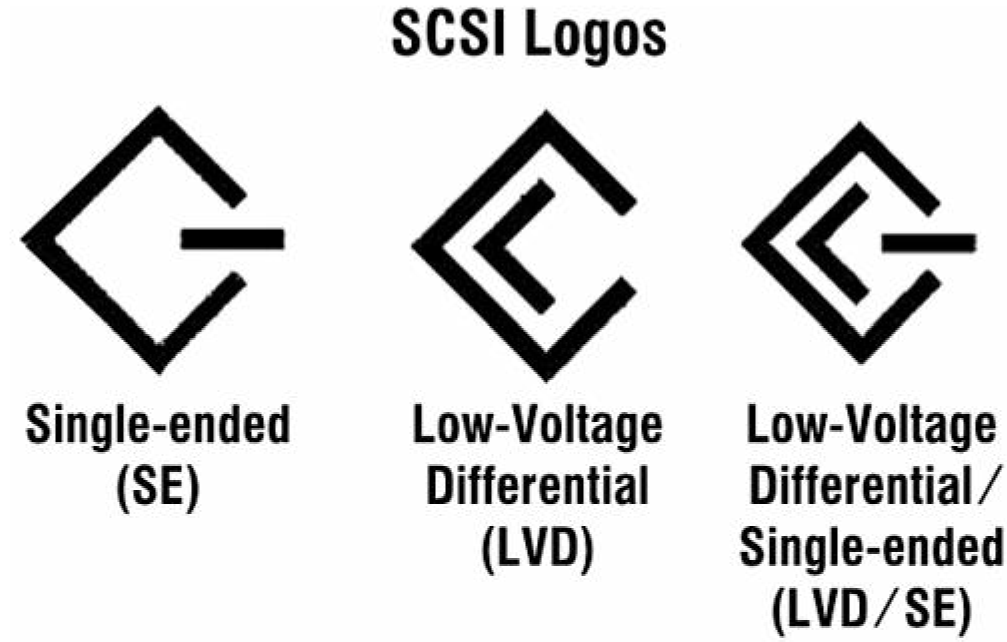

- Open the Device Manager display of onboard hardware to see if a SCSI device is listed; note that the Promise Ultra IDE controller built into some motherboards is listed under SCSI Devices in Device Manager, but is used for IDE/ATA devices. A true SCSI card will list SCSI in the device description in Device Manager.

Most SCSI cards plug into the PCI slot on the motherboard or the PC Card slot found in notebook computers. In these cases, the IRQ used by the card can be shared with other devices. However, older SCSI cards used ISA, EISA, or VL-Bus slots (EISA and VL-Bus were based on ISA); EISA IRQs can be shared with other EISA cards only, while ISA and VL-Bus cards can't share IRQs at all. Narrow SCSI (25-pin or 50-pin interface) cards support up to seven unique devices plus the host adapter, while Wide SCSI (68-pin) cards support up to fifteen unique devices plus the host adapter. External SCSI devices have two ports, enabling you to create a daisy-chain of devices. Each device plugged into a SCSI card must have a unique device ID, and the end of the daisy-chain of devices must be terminated. For more details about SCSI interfaces, daisy-chaining, and device configuration, see Chapter 14 , "Storage."

The USB Interface

The USB interface (refer to Figure 2.4) is replacing serial, parallel, and PS/2 mouse and keyboard ports. Some legacy-free systems feature only USB ports for external expansion, while others still have PS/2 keyboard ports, but leave other types of expansion up to the USB port. The original version of the USB port, USB 1.1, has speeds up to 12 megabits per second (Mbps), while the newer USB 2.0 standard, also called hi-speed USB, has speeds up to 480Mbps and is backward-compatible with USB 1.1. Most recent computers built through mid-2002 include built-in USB 1.1 ports, but computers built from mid-2002 on might include both USB 1.1 and USB 2.0 ports. Both USB 1.1 and USB 2.0 ports can be added to systems through the use of a PCI or PC Card add-on card.

USB ports can be used for keyboards, mice and pointing devices, scanners, printers, removable-media drives, hard drives, optical drives, direct data transfer, modems, and networking. USB ports are thus the most versatile ports built into modern PCs.

USB devices can be daisy-chained, as can SCSI and IEEE-1394a devices, but there are several differences between how these technologies work, as shown in Table 2.1.

Table 2.1. SCSI, IEEE-1394a, and USB Compared

Interface	How Devices Are Daisy-Chained	Maximum Number of Devices	How Devices Are Configured	Maximum Speed
SCSI	Direct connection between devices	7 or 15 (depends upon host adapter type)	Unique device ID; last device in daisy-chain must be terminated	10Mbps up to 640Mbps; speeds beyond 40Mbps are used primarily for hard drives
USB	Multiport hubs	Up to 127	Plug-and-play configuration in Windows	1.1: 12Mbps 2.0: 480Mbps
IEEE-1394a	Multiport hubs or direct connection between devices	Up to 16 with daisy-chaining; up to 63 through the use of hubs	Plug-and-play configuration in Windows	400Mbps

The USB ports built into the computer are known as root hubs; most computers have at least two. External hubs, also called generic hubs, are connected to root hubs to allow multiple USB devices to share a single root hub. Different USB devices use different amounts of power. Bus-powered hubs (hubs that take power from the USB root hub) can provide no more than 100 milliamps (mA) of power to each device, while self-powered hubs can provide the full 500mA of power required by some USB devices. Because USB root hubs, whether built into the motherboard or an add-on card, are PCI devices, they can share IRQs with other PCI devices.

USB 1.1 devices can run at 1.5Mbps (keyboards, mice, and pointing devices) or at speeds of 12Mbps (scanners, printers, and other devices). 1.5Mbps is referred to as low-speed USB 1.1, while 12Mbps is referred to as full-speed USB 1.1. Because low-speed devices can slow down full-speed USB 1.1 devices when they are plugged into the same USB 1.1 root hub, I recommend that you use separate root hubs for low-speed and full-speed USB 1.1 devices. However, if you have USB 2.0 ports (also called hi-speed USB ports), USB 2.0 is capable of managing both USB 1.1 and USB 2.0 devices at their top speeds. For more information about USB ports, see Chapter 8 , "Input/Output Devices and Cables."

The IEEE-1394a Interface

The IEEE-1394a interface (see Figure 2.8), commonly referred to as IEEE-1394, FireWire, FireWire 400, or i.Link, is designed to be a high-speed replacement for SCSI, parallel, and other **legacy port** types.

Figure 2.8. SCSI (left) and IEEE-1394a (right) host adapters installed in a typical PC.

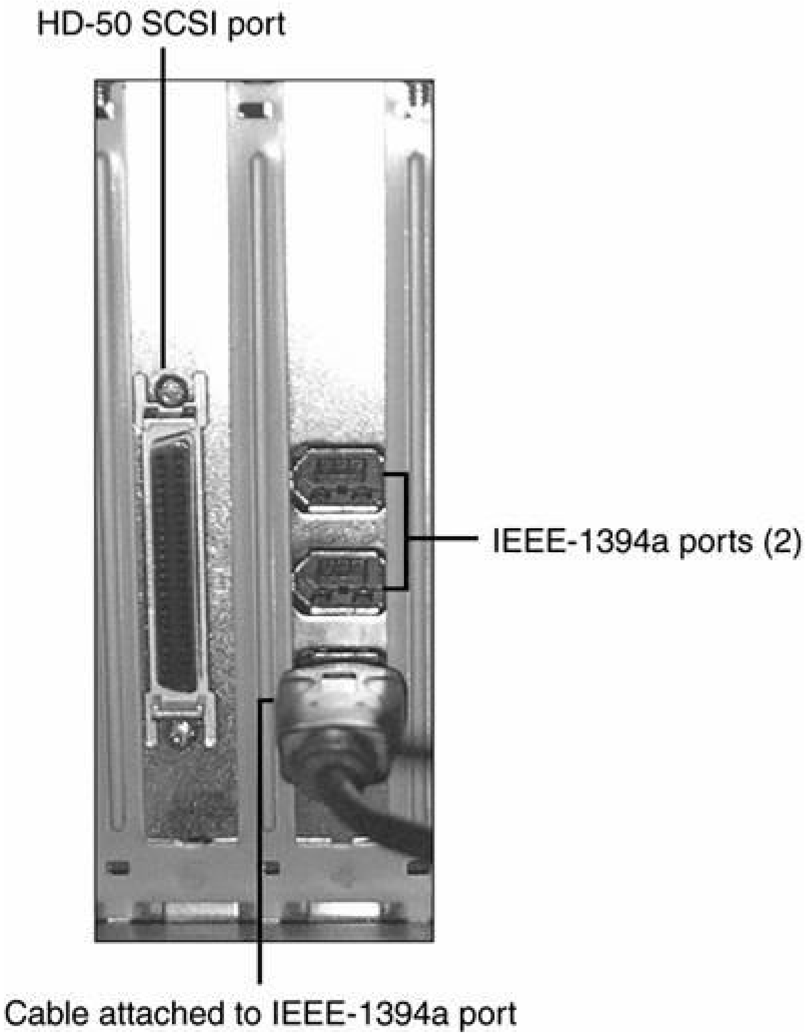

As Table 2.1 indicates, IEEE-1394a is much faster than USB 1.1, but a bit slower than USB 2.0. The major differences between IEEE-1394a and USB, aside from speed, is that IEEE-1394a can be used to connect devices without using a computer and provides enough power to run drives and scanners (when used with a six-wire cable), whereas USB 1.1 and USB 2.0 require that a computer be used to control the connection and can't provide sufficient power for some types of devices. For these reasons, IEEE-1394a is very popular for use with DV camcorders and other multimedia devices. The speed of IEEE-1394a and its increasing popularity has also made it a popular choice for external hard drives. IEEE-1394a ports are occasionally built into a computer's motherboard, but are more often installed into PCI slots or CardBus PC Card slots. For more information about IEEE-1394a ports, see Chapter 8 .

Points of Failure on the Rear of Your Computer

The most likely point of failure on the rear of your computer is peripheral cabling. Fortunately, more and more devices use the lightweight USB cable shown in Figure 2.4 instead of the bulky, heavy serial and parallel cables also shown in Figure 2.4. Note that serial, parallel, and VGA cables all use thumbscrews; if you don't fasten the thumbscrews to the connector on the computer, the cables won't connect tightly, which can cause

intermittent or complete peripheral failure. Most SCSI cables use locking clips to hold them in place. If the clips aren't engaged, this can cause intermittent or complete SCSI peripheral failure.

Newer types of peripheral cables such as USB and IEEE-1394a are pushed into place and are very lightweight. No thumbscrews or other locking devices are needed. However, these cables can also be pulled out of the socket easily, precisely because they are lightweight and support a feature called hot-swapping. Hot-swap devices can be freely connected and disconnected while the PC's power is on.

caution

When you attach cables to the ports at the rear of the computer, avoid tangling them together. Tangled cables can cause electrical interference with each other, leading to erratic performance of external devices such as your printer or monitor. Also, tangled cables put extra stress on ports, which can cause malfunctions or port failure.The power supply shown in Figure 2.5 is another likely point of failure. If the three-prong power cable is not plugged all the way into the computer, the system might not start up at all, or might shut down unexpectedly. If the voltage selector switch is not set correctly, the computer will not start at all, and if the power supply is set for 115V and is plugged into a 230V supply, the power supply and possibly other parts of the computer will be destroyed.

All Around a Laptop Computer

Notebook computers use the same types of peripherals, operating system, and application software as desktop computers use. However, notebook computers vary in several ways from desktop computers:

- Most notebook computers feature integrated ports similar to those found in recent desktop computers (such as USB 1.1 or 2.0 ports and 10/100 Ethernet network ports) as well as one or more PC Card (PCMCIA) slots and a 56 kilobits per second (Kbps) modem.
- Some notebook computers support swappable drives, but less-expensive models require a trip to the service bench for a drive upgrade.
- More and more notebook computers use combo DVD-ROM/CD-RW drives to enable one optical drive to perform the work of two.
- Many notebook computers don't have an internal floppy drive, but rely on CD-RW or removable-media drives on USB connections to transfer or back up data.
- Notebook computers have integrated pointing devices built into their keyboards; most use a touchpad, but a few (primarily IBM and some Toshiba models) have a pointing stick (which is better is a matter of personal preference).

Figure 2.9 shows you a composite view of a typical notebook computer, a Compaq Presario 700-series.

Figure 2.9. Rear, left, and right views of the author's Compaq Presario 700-series notebook computer.

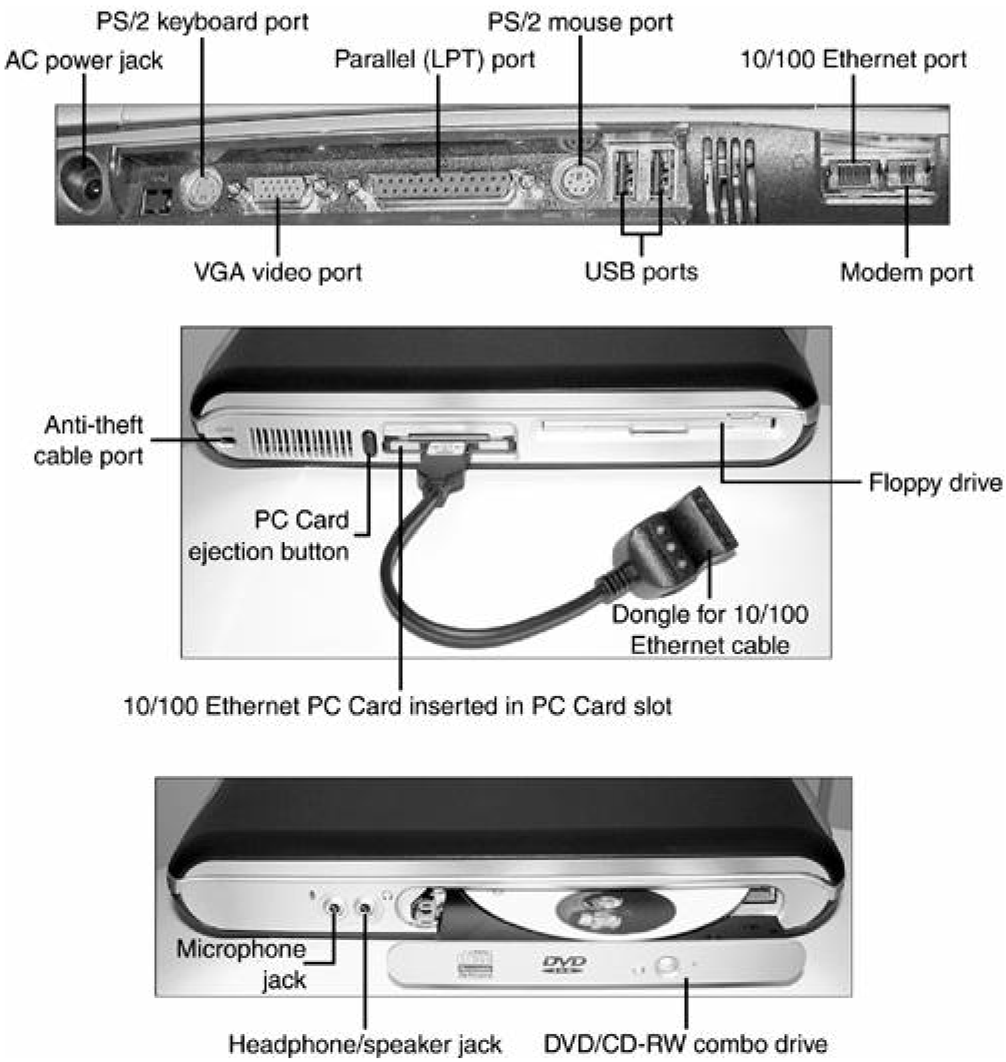

Points of Failure on a Notebook Computer

As with desktop computers, cabling can be a major point of failure on notebook computers. However, notebook computers also have a few unique points of failure. The PC Card (PCMCIA card) represents a significant potential point of failure for the following reasons:

- If a PC Card is not completely pushed into its slot, it will not function.
- If a PC Card is ejected without being stopped by using the PC Card system tray control, it could be damaged.
- Many PC Cards designed as 10/100 Ethernet network adapters or 56Kbps modems use dongles similar to the one pictured in Figure 2.9. If the dongle is damaged, the card is useless until a replacement dongle is obtained.
- Some notebook computers that have the 32-bit CardBus version of the PC Card slot require the user to enable CardBus compatibility in the system BIOS, or else CardBus cards (used by USB 2.0, IEEE-1394a, and other high-bandwidth PC Card devices) will not work.

For more information on PC Card and CardBus devices, see Chapter 12 , "Portables."

Although a notebook computer's drives are much more rugged than those found in desktop computers, they are much more expensive to replace if damaged. Although some mid-range and high-end notebook computers offer swappable drives, most lower-priced models do not. You can perform an upgrade to a hard disk without special tools on many models, but replacement of other types of drives on systems that don't support swappable drive bays can be expensive.

Although drives are expensive to replace on notebook computers, the biggest potential expense is the LCD screen. Most recent computers use active-matrix LCD panels in place of the dimmer (but less expensive to fix) passive-matrix panels once common.

Inside a Typical PC

As you have already learned, some problems that manifest themselves on the outside of the computer come from problems inside the computer. If you ever add memory, add an internal drive, upgrade your processor or motherboard, or add a card to your computer, you will need to work with the interior of the computer to complete these tasks.

The interior of a typical desktop computer is a crowded place, as Figure 2.10 shows.

Figure 2.10. The interior of a typical PC using an ATX motherboard.

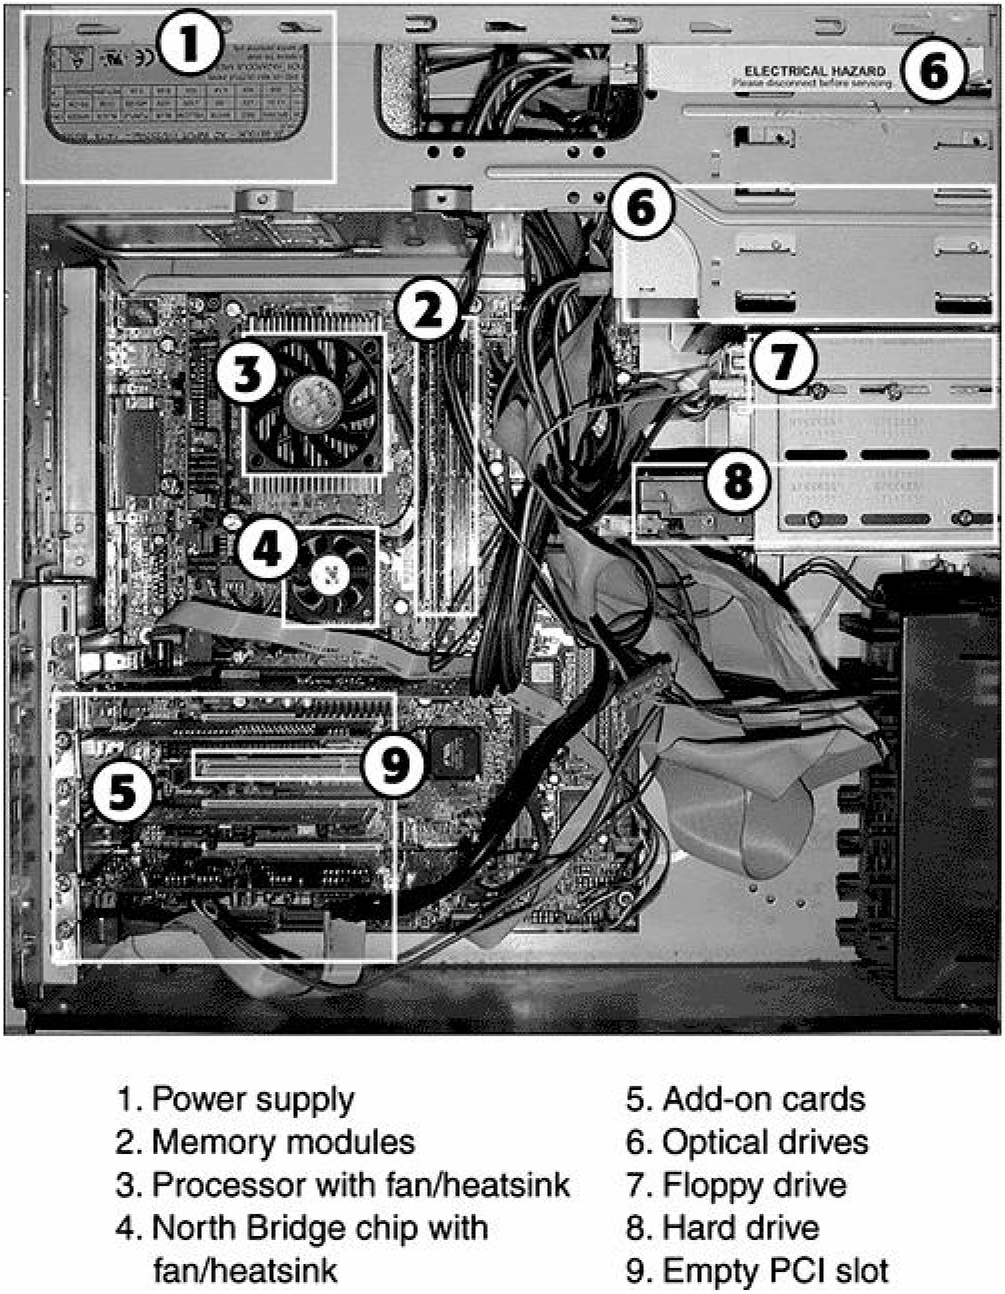

Each of the devices highlighted in Figure 2.10, as well as the data, signal, and power cables that connect them to the motherboard and power supply, can cause significant system problems if they fail.

caution

If you're not familiar with working inside your PC, read "Working Inside Your PC," later in this chapter, for help in opening the case and preventing damage to components inside your PC.

Expansion Slots

Although today's systems have more built-in ports than ever, sooner or later you might want (or need) to add a new port or card type. That's a job for an expansion slot.

Typical mini-tower and full-tower desktop computers have three or more expansion slots, some of which might already be used for factory-installed devices such as video cards, network cards, or modems. Most computers have several PCI slots, and many also have a single AGP slot for high-speed video. Although AGP slots are faster than PCI slots, they are configured the same way in the system BIOS.

See "PCI Configuration" in Chapter 6 for details.

Most recent systems have no more than one ISA slot, and many no longer have ISA slots at all. ISA cards are much slower than PCI cards, and some can't be installed with Windows's normal plug-and-play automatic detection. You should avoid using ISA slots for upgrading your system if possible. Figure 2.11 compares AGP, PCI, and ISA slots.

Figure 2.11. ISA, PCI, and AGP slots compared.

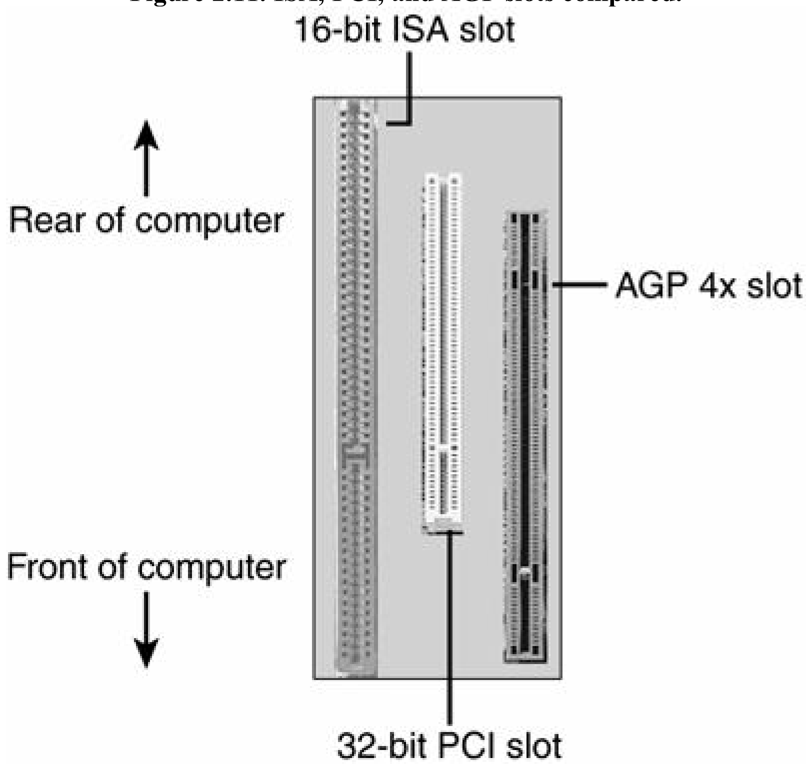

Regardless of the type of expansion slot an add-on card uses, you need to push the card connector all the way into the expansion slot when you install a card, as shown in Figure 2.12.

Figure 2.12. A video card partly inserted into the slot (top) and fully inserted into the slot (bottom).

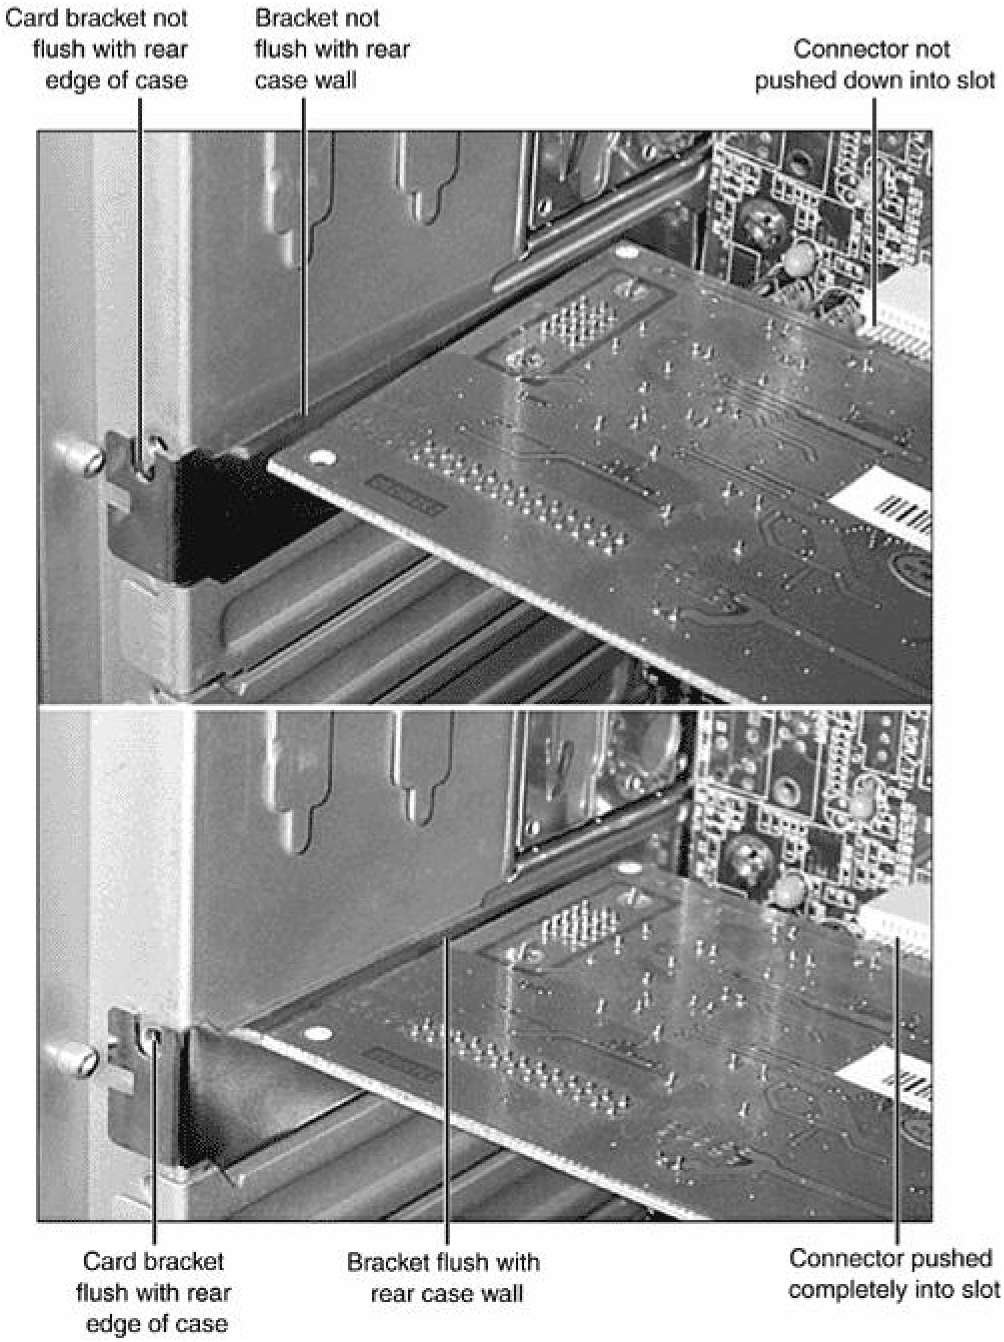

After the card is properly inserted into the expansion slot, you need to fasten the card to the case with a screw.

Points of Failure Inside the Computer

Some of the problems you could encounter because of devices inside your computer include

- Overheating— Failure of the fans in the power supply or those attached to the processor, North Bridge chip, or video card can cause overheating and can lead to component damage. Each fan shown in Figure 2.10 is connected to the

motherboard to obtain power. Some case-mounted or older processor fans use a standard four-wire drive power connector instead (shown later in Figure 2.15).

Figure 2.15. A typical ATA/IDE CD-ROM drive before (top) and after (bottom) typical cabling and configuration problems have been corrected.

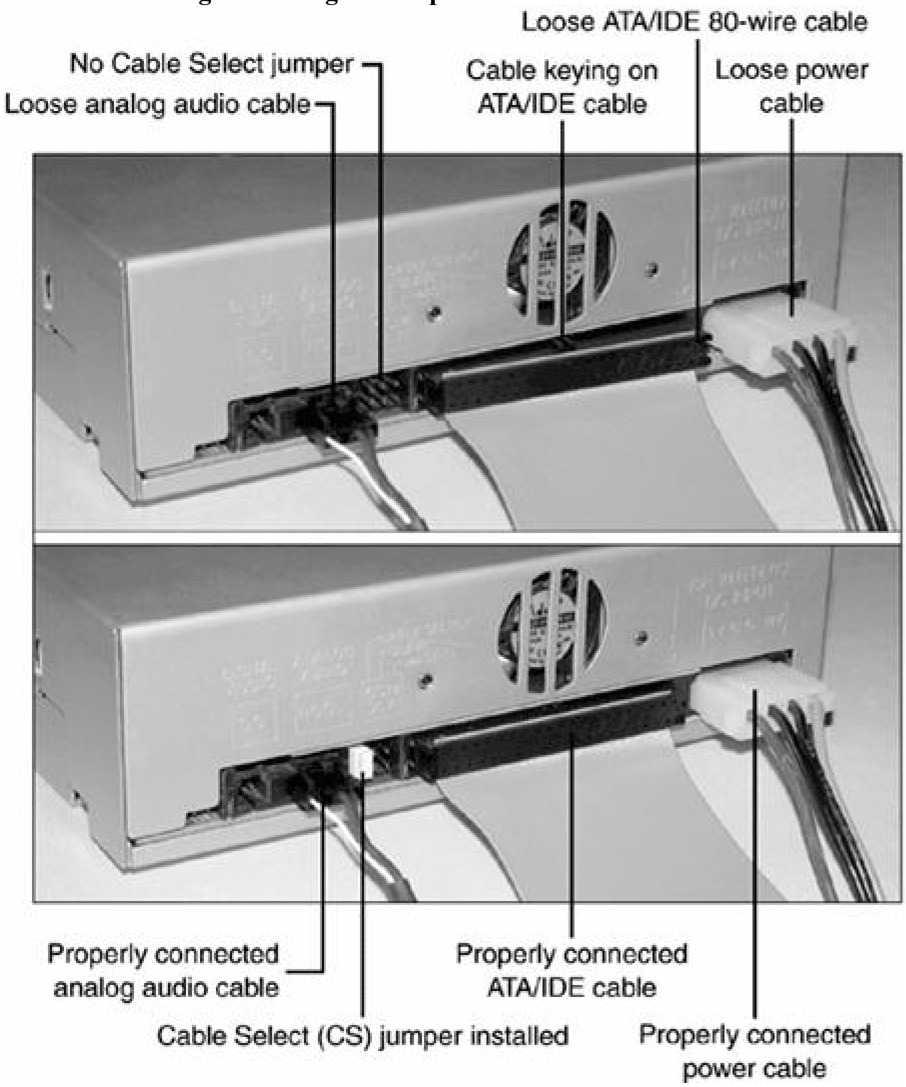

- Loose add-on cards (see Figure 2.12)— A loose add-on card might not be detected by plug-and-play or the Windows Add New Hardware Wizard, or might have intermittent failures after installation.
- Inability to start the computer— A loose processor or memory module can prevent the computer from starting (see Figures 2.13 and 2.14).

Figure 2.13. A socket-based processor before (left) and after (right) serious problems were corrected.

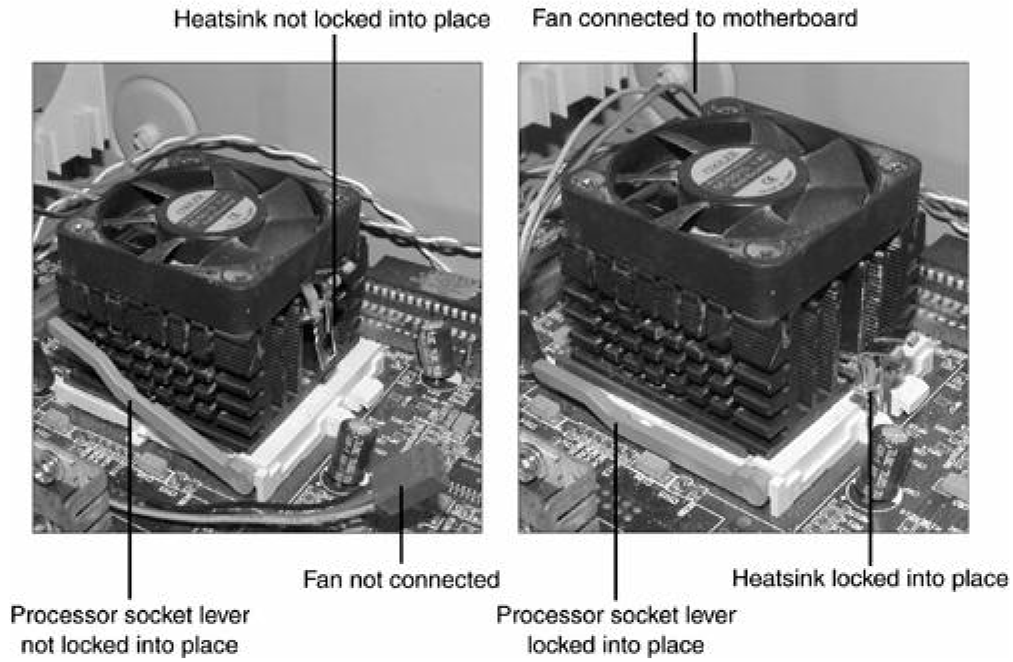

Figure 2.14. A memory module before (top) and after (bottom) being locked into its socket.

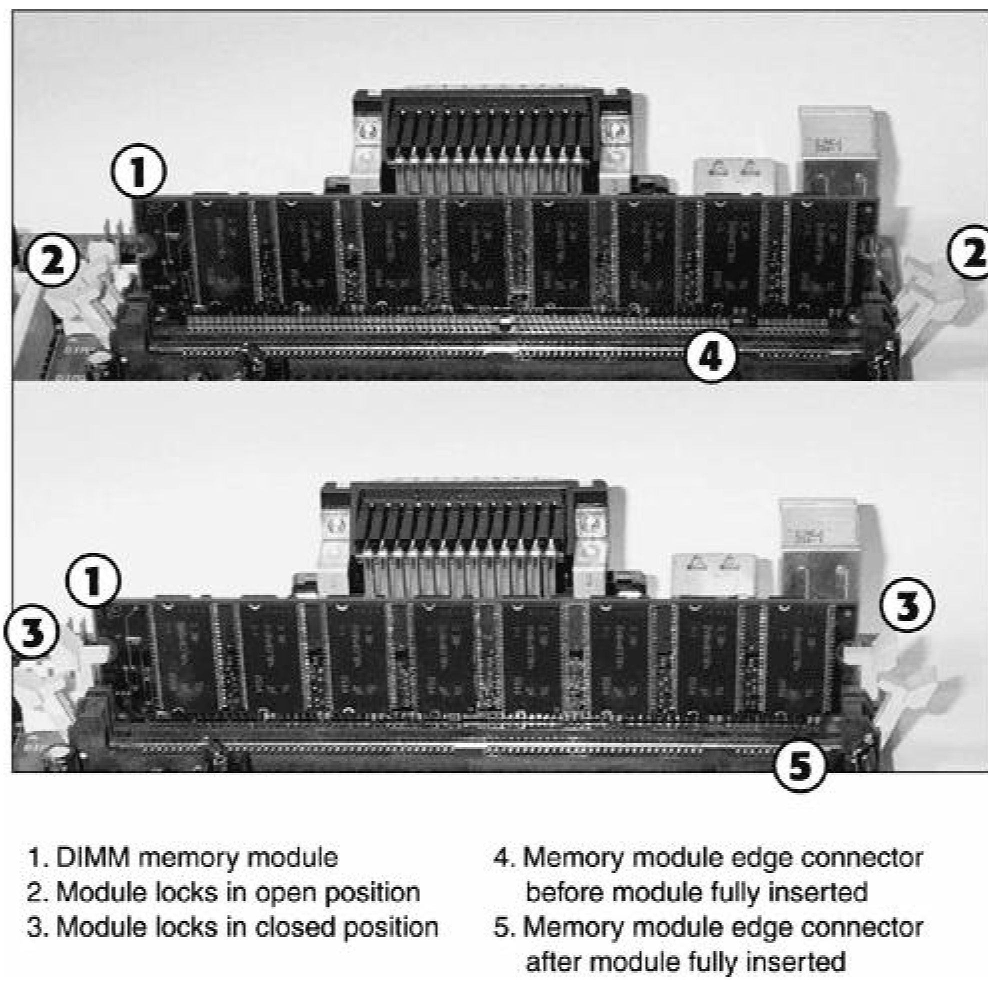

- Drive failures— If drives are not properly connected to power or data cables, or are not properly configured with jumper blocks, they will not work properly (see Figures 2.15 and 2.16).

Figure 2.16. A typical ATA/IDE motherboard drive connector (top) and the DIP switches used on some motherboards for system configuration (left).

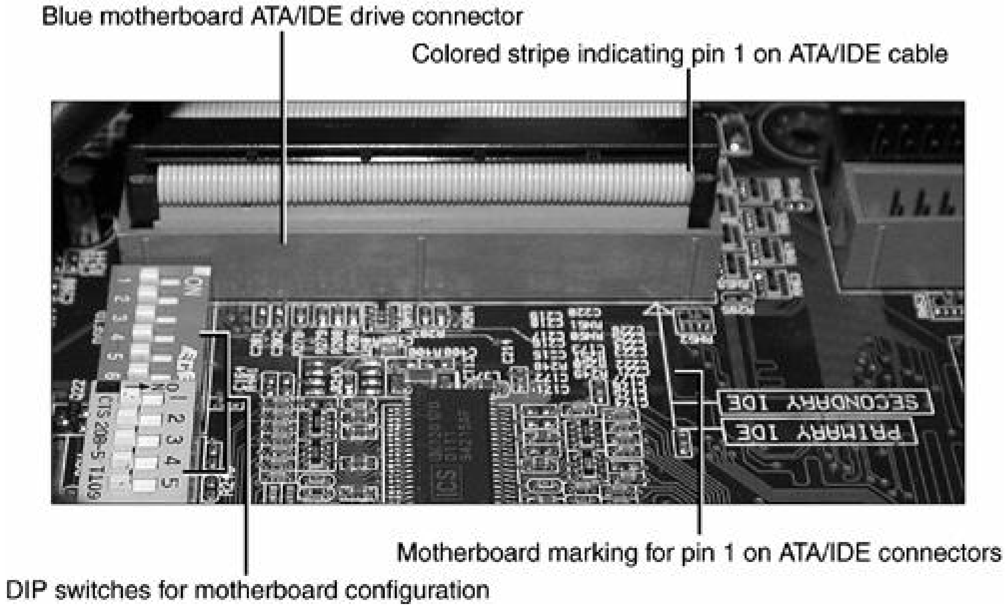

- Multimedia failures— If the analog or digital audio cable running from the CD-ROM or other optical drive to the sound card is disconnected, you might not be able to hear CD music through the speakers. Some high-end sound cards also have connections to an external breakout box for additional speaker or I/O options. Be careful when you work inside a computer to avoid disconnecting these cables (see Figure 2.17).

Figure 2.17. The top edge of a typical high-end sound card (a Sound Blaster Audigy Platinum) with multiple cable connections.

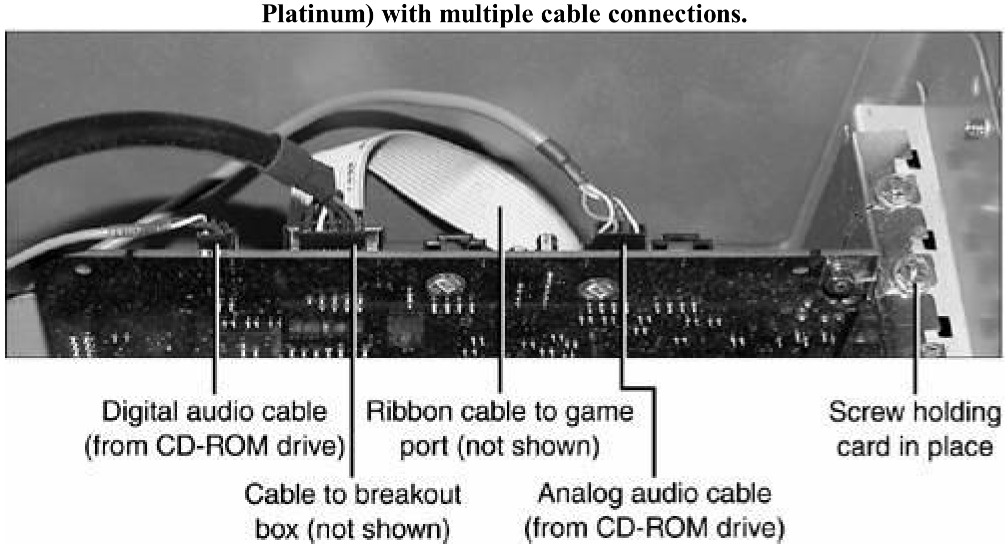

- Front panel failures— The tiny cables that connect the case power switch, reset switch, and status lights are easy to disconnect accidentally if you are working near the edges of the motherboard (see Figure 2.23).

Figure 2.23. The front-panel cables, BIOS chip, and battery on a typical motherboard.

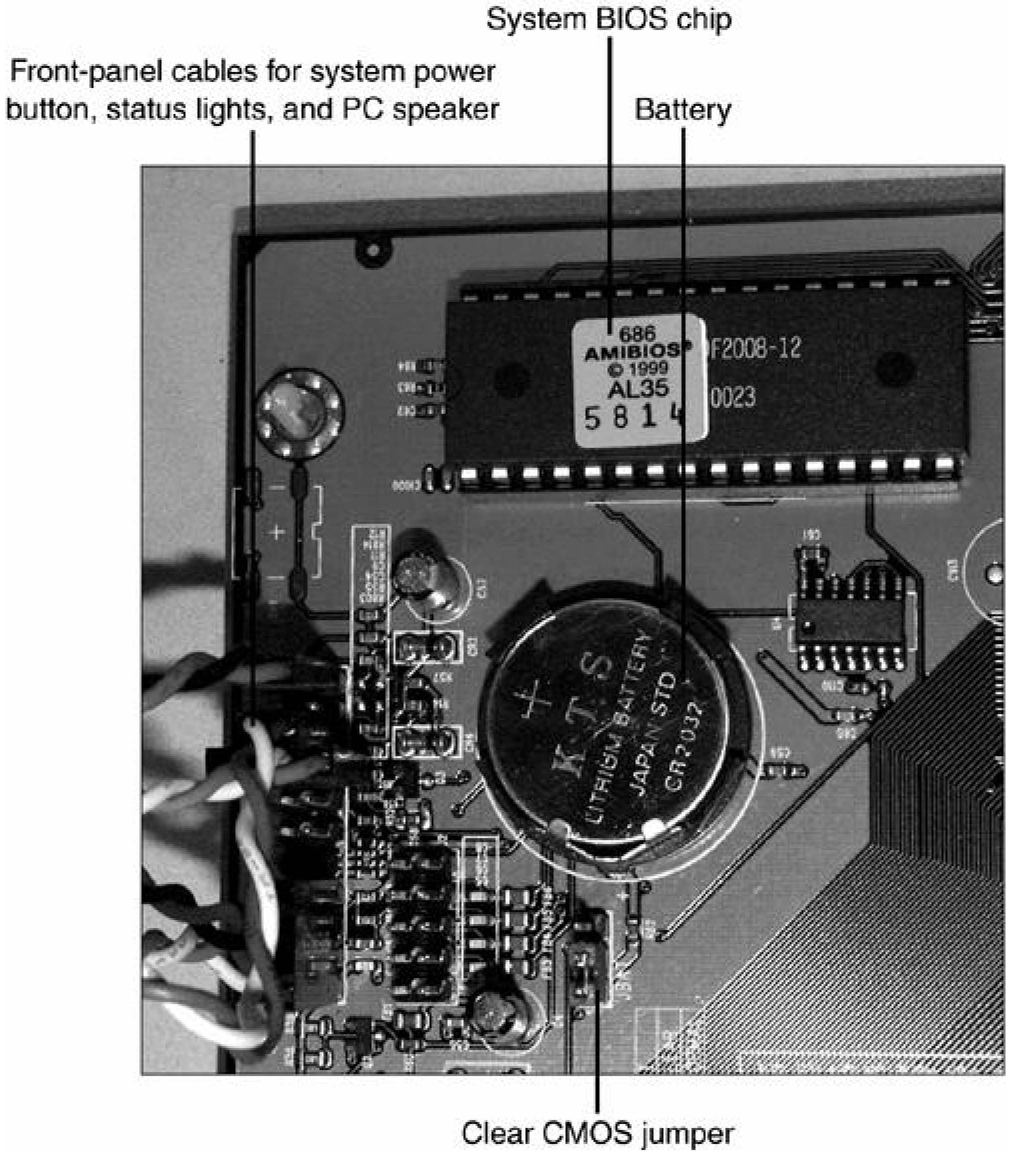

- Battery failure— The battery (see Figure 2.23) maintains the system settings that are configured by the system BIOS. The settings are stored in a part of the computer called the CMOS, more formally known as the non-volatile RAM/real-time clock (NVRAM/RTC). If the battery dies (the average life is about two to three years), these settings will be lost. After you replace the battery, you must re-enter the CMOS settings and save the changes to the CMOS before you can use the system.
- BIOS chip failure— The system BIOS chip (see Figure 2.23) can be destroyed by electro-static discharge (ESD) or lightning strikes. However, BIOS chips can also become outdated. Although some systems use a rectangular socketed BIOS chip like the one shown in Figure 2.23, others use a square BIOS chip that might be socketed or surface-mounted. In both cases, software BIOS upgrades are usually available to provide additional BIOS features, such as support for newer processors and hardware.

The following sections introduce you to these components and how they work.

The ATA/IDE Interface

Virtually every desktop computer built since the mid-1990s has featured two ATA/IDE interfaces on the motherboard. Each ATA/IDE interface can handle one or two drives, including hard drives, optical drives, and removable-media drives. Thus, you can install up to four ATA/IDE drives into a typical desktop computer.

tip

Want to install more ATA/IDE drives? Add a PCI-based ATA/IDE host adapter. Some of these cards also support ATA RAID or Serial ATA, which can provide greater speed, data security, or easier installation. See Chapter 14 for more information about ATA RAID and Serial ATA.The ATA/IDE interface uses a 40-pin connector (see Figure 2.15), which connects to a 40-wire or 80-wire ATA/IDE cable (both cable types use the same connector). This cable has three connectors, which go to the following locations:

- ATA/IDE interface on the motherboard
- Master drive
- Slave drive

40-wire cables, which are now becoming outdated, support device speeds up to UltraDMA/33 (33 megahertz [MHz]). 80-wire cables support all older devices plus the latest drive speeds up to 133MHz (UltraDMA/133). Because most ATA/IDE drives on the market today run at UltraDMA/100 or UltraDMA/133 speeds, you should not use 40-wire cables with hard disk drives anymore.

caution

Most ATA/IDE drives and motherboard host adapters are designed to accept cables only when they are properly connected. The usual method is to match a key on one side of the cable plug (refer to Figure 2.15) with a matching cutout on the cable connector. However, some low-cost cables don't feature keying, making it easy to install them incorrectly. If you use such a cable, note that the colored stripe indicating pin 1 on the cable is usually on the same side of the cable as the power supply connector. Check the motherboard (shown later in Figure 2.16) for markings.Each ATA/IDE interface uses an IRQ: IRQ 14 is used for the primary interface, and IRQ 15 is used for the secondary interface. These IRQs cannot be shared with other devices. For details, see "Hardware Resources ," later in this chapter.

Even if you aren't installing a new drive, if you move existing ATA/IDE cables around inside your computer to gain access to memory, processor, or other components, you need to recheck the cable connections, both to the drive and to the motherboard. It's very easy to accidentally pull these cables loose. Whether a cable is completely removed or only partially detached, drives do not function properly.

caution

The original ATA/IDE drive cable, which contained 40 wires, enabled users to select primary (master) and secondary (slave) drives with jumper blocks on the rear or bottom of the drives. Some vendors, such as Western Digital, supported a no-jumper-block

configuration for single hard drives. Most drives are labeled with the correct jumper settings as well as with the correct orientation for the power and data cables. If your drive lacks this information, look up the drive model on the vendor's Web site. However, in most new systems today, ATA/IDE hard drives are connected with an 80-wire ATA/IDE cable. This enables ATA/IDE drives such as hard drives, CD-RW drives, or DVD drives to be jumpered as Cable Select, giving control of master/slave settings to the cable. Most of these cables are color coded. With color-coded cables, the black connector at one end of the cable is used for master, the middle (gray) connector is used for slave, and the blue end of the cable connects to the system board or other ATA/IDE host adapter.Figure 2.16 shows a typical ATA/IDE motherboard connector.

Audio Hardware

Although more and more motherboards feature built-in sound, serious gamers and hardcore audio fans still insist on a separate sound card. Many high-end sound cards, such as the Creative Labs Sound Blaster Audigy shown in Figure 2.17, can be connected to breakout boxes that provide additional audio connectors and controls, and some also feature onboard IEEE-1394a ports.

Both motherboard-integrated audio and separate sound cards can also be connected to CD-ROM and other optical drives with four-wire analog or two-wire digital audio cables, also shown in Figure 2.17.

Working Inside Your PC

If you've never opened up a computer before, it can be pretty overwhelming. In this section, I'm going to help you get started with practical advice on how to

- Open the case.
- Protect your system against electrostatic discharge (ESD).
- Connect internal and external data cables.
- Install a PCI card.

Some systems can transfer CD music through the standard ATA/IDE cable and don't require a separate patch cable. If you can play music CDs through your speakers and your system doesn't use analog or digital audio cables, don't worry about it. Most CD-ROM and other types of optical drives include analog or digital audio cables, but if you need replacements, most computer stores also stock them.

Opening the Case of a Desktop PC

I recommend that you look at your system manual for case-opening instructions, particularly if you have a retail-store system made by HP or Compaq. Depending upon the type of case you have, you might need to remove just one or two screws, or maybe a handful. If you're opening the case to gain access to the motherboard, you might need to do more than just take the cover off the system.

So-called "white box" systems are usually pretty straightforward to open, because they use case designs made for user access instead of low cost. Figure 2.18 shows the rear of two typical cases used by white-box computer dealers or as replacement cases. The computer on the left has a single screw holding the covers in place. After this screw is removed, the top panel must be removed before the side panels can be removed. The computer on the right uses four screws per side to hold the side panels in place, but the side panels can be

removed without removing the top cover. The right-side panel can even be swung out and latched back into place for faster card and drive installation.

Figure 2.18. Two generic cases compared to each other.

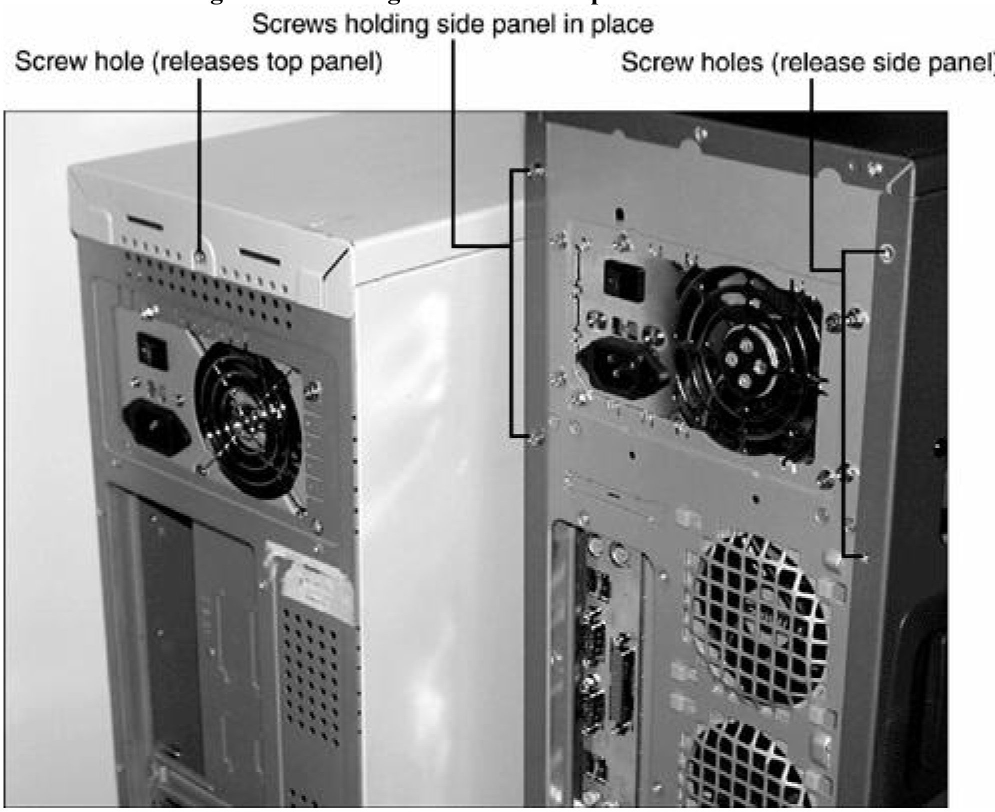

Retail-store systems often use a single-piece molding that slips off the rear of the chassis. It can be held in place by several screws.

Taking ESD Precautions

After you open your PC, what should you do next?

caution

Be sure to unplug the system before you open the case. Don't mix up the screws used for the case with screws used for holding expansion cards or drives in place. Keep them separate. Label empty film or medicine containers and use one for each type of screw you remove.Ideally, you should wear a commercial wrist strap made for ESD protection and clip it to an unpainted metal part on the computer you are servicing, as in Figure 2.19, before you touch any other part of the interior of your computer.

Figure 2.19. A wrist strap and some suitable places to clip it inside the system.

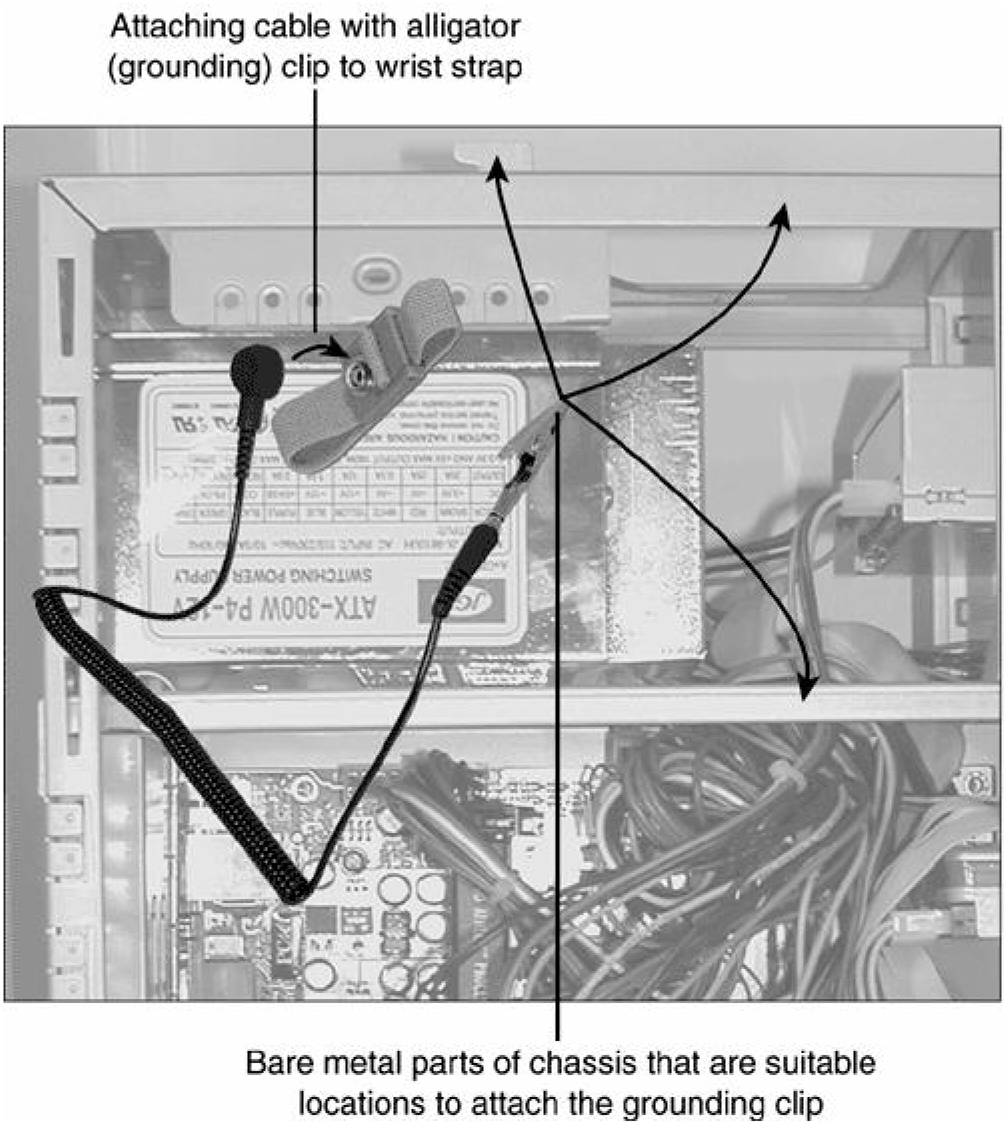

When you put the wrist strap on, make sure the metal plate on the inside of the strap touches bare skin; don't wear it over a shirt or sweater, and make sure it's comfortably snug. When the wrist strap is connected to a metal part of the chassis with the alligator clip on the end of the cable that snaps onto the wrist strap, it equalizes the electrical potential of your body and the computer to prevent ESD discharge.

However, if you don't have a wrist strap, touch unpainted metal on the chassis or the power supply before you touch or remove any cables or other parts. If possible, keep touching the chassis or power supply while you work.

For more details on ESD protection, see Chapter 13, "Safety and Recycling."

Connecting Internal and External Data Cables

Connecting internal and external drives and devices is critical if you want to have reliably working PCs. Here are a couple of examples of how to do it right:

Attaching Cables to the Floppy Drive and Controller

The floppy drive uses a 34-pin cable that has a twist at one end. This end of the cable connects to the floppy drive (A: drive). The middle connector is used for the B: drive (if present; not supported on some machines). The connector with the untwisted end plugs into the floppy controller on the motherboard. The colored marking on the cable and the twist indicate the pin 1 side of the cable. This is important to note because floppy drives and controllers are not always keyed to prevent incorrect installation.

The floppy drive is powered by a small 4-pin power cable. There is a small cut-out in the center of the connector on the drive that corresponds to a projection on the cable. Figure 2.20 shows how the floppy drive cable and power cable connect to the rear of a typical floppy drive and to the floppy controller.

Figure 2.20. Connecting data and power cables to a typical 3.5-inch floppy drive (top) and floppy controller on motherboard (bottom).

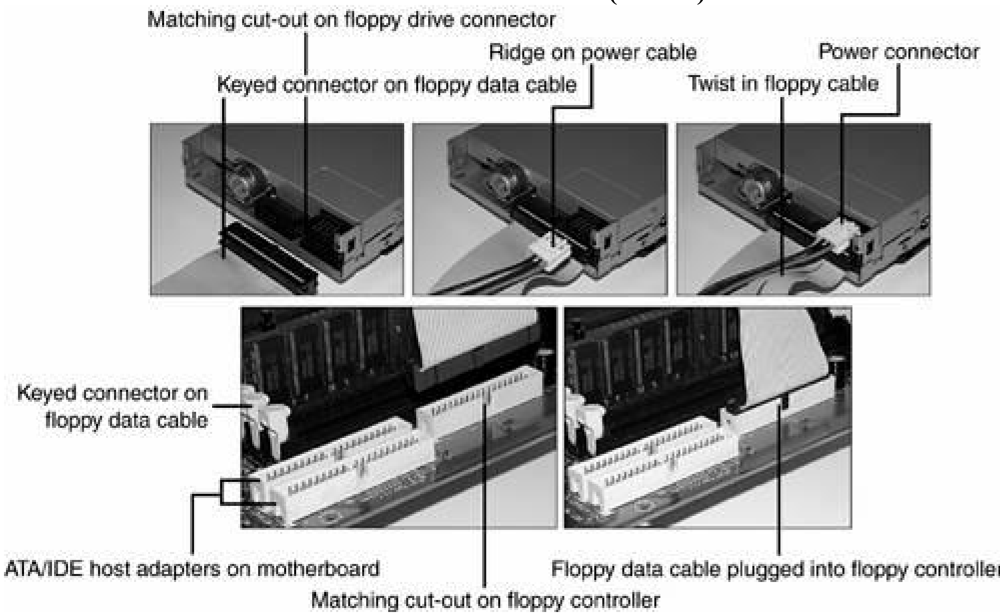

To install the cables

- Take the ESD precautions discussed in "Taking ESD Precautions," earlier in this chapter, after you open the case and before you touch the cables or other components.
- Be sure to line up the keying on the cable (if any) with the cut-out on the floppy drive and controller connectors.
- Push the cable firmly but gently into place.
- Make sure the ridge on the power cable connector faces away from the drive; the power cable can be forced into place upside down, but this will destroy the drive when the power is turned on.
- Push the power cable connector firmly but gently into place. If you don't have a spare 4-wire connector for the floppy drive, you can purchase an adapter, which converts the large Molex connector used for hard disks into a floppy drive power connector. You can also purchase a Y-splitter, which can power a floppy drive and a hard drive.

caution

If any pins are bent on the floppy drive, controller data, or power connectors, carefully straighten them with a pair of needlenose pliers before installing the cables.The installation of an ATA/IDE hard disk or other drive type is similar in many ways, but the cables used are different and the drives must be correctly jumpered.
For more information on floppy drives and hard drives, cables, and floppy media, see Chapter 14 .

Attaching the VGA Cable to a Video Card or Port

The VGA cable has as many as 15 wires that are routed into a connector the same size as a serial (COM) port. Consequently, it is one of the heaviest cables you need to connect to a PC. If you don't attach it correctly, you could have poor-quality images on your monitor, and if the cable falls off, you won't have any picture at all.
Depending upon the computer, the VGA port might be located on a card built into an expansion slot, or might be clustered with other ports at the rear of the computer. In either case, the port and the cable are the same.
Figure 2.21 shows how the VGA cable should be connected to the VGA port.

Figure 2.21. Connecting a VGA cable to a VGA port.

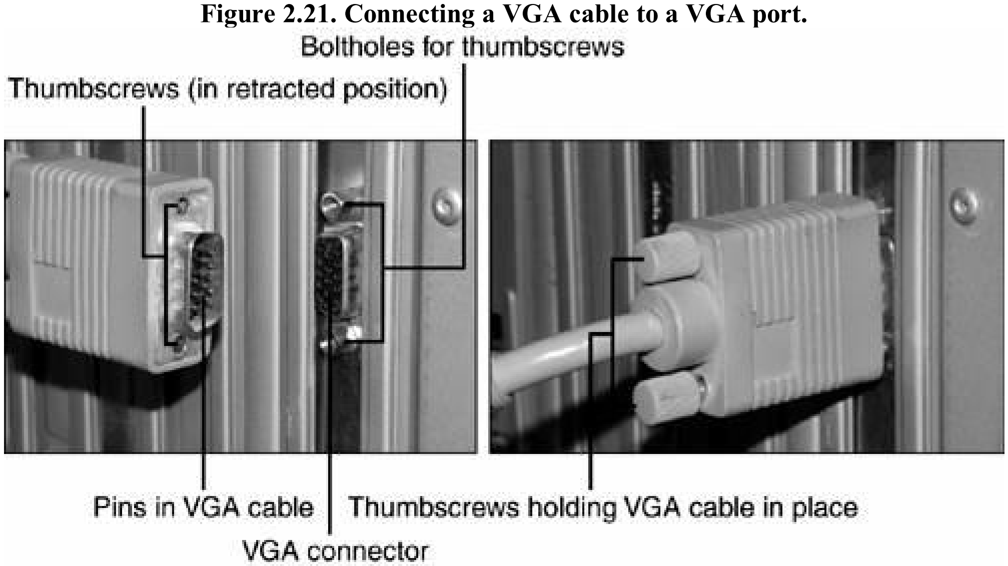

To connect the VGA cable to the VGA port

- Turn the thumbscrews so they are completely retracted. If they are not completely retracted, they could prevent proper connection.
- Check the VGA cable for bent pins, but don't panic if you see a few pins that appear to be missing. Most VGA monitors don't use all 15 pins. Carefully straighten any bent pins you find.
- Line up the cable with the port and carefully push the cable into the port.
- Fasten the thumbscrews tightly but evenly to hold the cable in place.

For more information about graphics cards and cables, including installing and configuring cards that support multiple displays, see Chapter 9 , "Video."

Installing a PCI Card

Most systems have at least one open PCI card slot; it's the standard card type used for most devices. If you need to add an internal modem, a network adapter, a USB 2.0, an IEEE-1394a, a hard disk, or a SCSI host adapter to a desktop PC, you will add a PCI card that includes the necessary port.

Figure 2.22 shows the process of adding a PCI card (specifically, a Serial ATA hard disk host adapter) to a typical system. The procedure for installing an AGP-based video card is similar, except that the AGP slot is used.

Figure 2.22. Installing a Serial ATA card into a PCI slot.

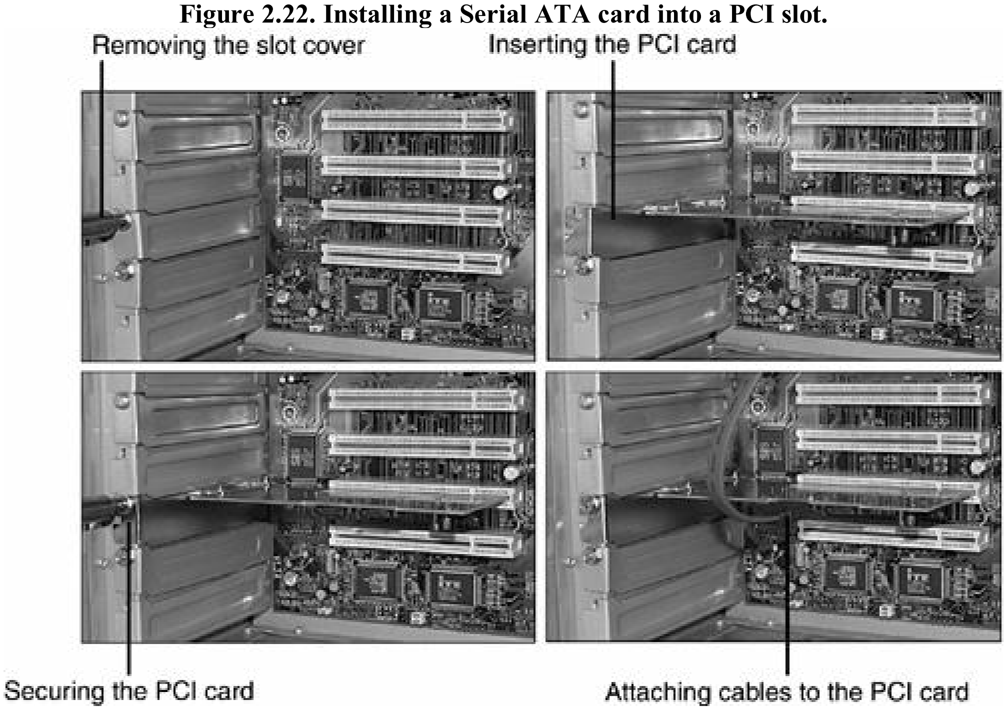

To add a PCI card

1. Shut down the system and unplug it.
2. Open the system.
3. Follow the ESD precautions discussed in "Taking ESD Precautions," earlier in this chapter.
4. Remove the slot cover behind the PCI slot you want to use.
5. Line up the connector on the bottom of the card with the slot.
6. Carefully push the card into place until it locks into place. Make sure the bracket on the card fits between the rear edge of the motherboard and the outer wall of the case.
7. Fasten the card into place using the screw you removed from the slot cover in step 4.
8. Attach any cables required between the card and other devices.
9. Double-check your work; then close up your system and follow the instructions provided with the card for driver installation and card configuration.

For more information about installing cards and configuring them, see Chapter 8 , "Input/Output Devices and Cables"; Chapter 9 , "Video"; Chapter 14 , "Storage"; Chapter 19 , "Installing and Configuring Hardware in Windows"; and Chapter 21 , "Networking and Internet Connectivity."

The System BIOS

The system BIOS chip shown in Figure 2.23 is responsible for configuring many parts of your computer, including

- Floppy, optical, and hard drive configuration
- Memory size and speed
- Drive boot sequence
- Built-in port configuration
- System security
- Power management
- Plug-and-play hardware configuration
- Processor compatibility and speed setting

Essentially, the BIOS acts as a restaurant menu of possible choices, and the CMOS RAM (which might be a separate chip or built into the South Bridge on some chipsets) stores the selections made from the menu of choices. When you received your computer from the factory, default selections were already stored in the BIOS, but as you add devices or customize your computer to perform certain operations, you might need to make additional choices. The battery shown in Figure 2.23 maintains the contents of CMOS memory; a typical battery lasts for two years or longer.

How the BIOS Displays Your PC's Components

Although some computers display only a system manufacturer's logo at startup, forcing you to read the system manual to determine which key to press to start the BIOS setup program, others, particularly "white box" computers that use a collection of components from various vendors, or systems that use a replacement motherboard, can provide you with a lot of useful information at startup.

Figure 2.24 shows a typical example of the BIOS chip's POST (power-on self-test) program detecting onboard storage (the memory size is displayed briefly on many systems first). The display also shows which key to press to start the setup program.

Figure 2.24. A typical startup screen displaying detected drives, chipset, and BIOS information.

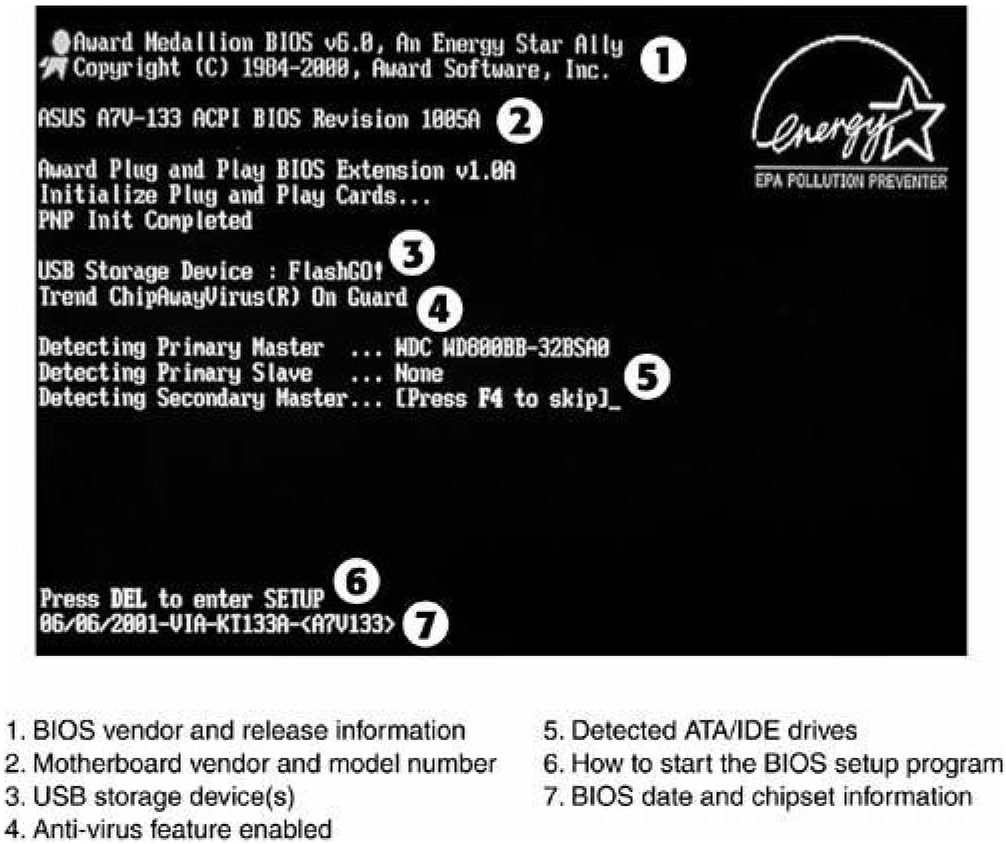

Many systems that display information similar to that shown in Figure 2.24 also display a condensed listing of onboard hardware before starting Windows (see Figure 2.25). Because this information should not change on a day-to-day basis unless you change your system configuration (BIOS changes or hardware upgrades), displaying this information at startup is a valuable aid to troubleshooting a sick system.

Figure 2.25. A typical system configuration screen displayed at system startup.

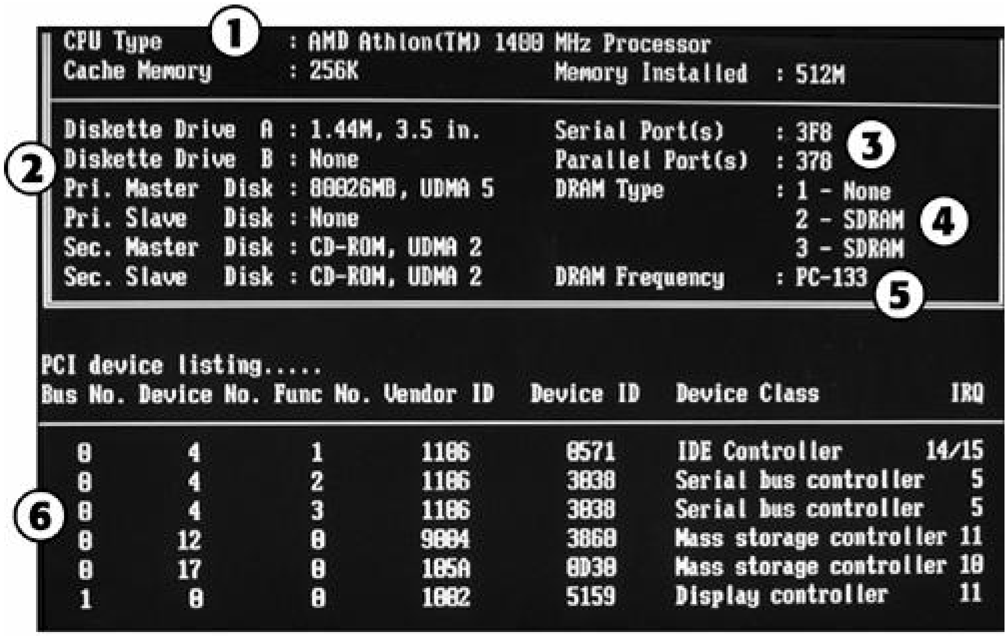

1. Processor type, speed, and memory size information
2. Drive information
3. I/O port addresses for serial and parallel ports
4. Memory slot usage
5. Memory speed
6. Onboard PCI devices

One of the reasons it's so important to display this information when you start your computer (if your system permits it) is because you will know immediately if there are any changes in your hardware configuration. If the configuration information displayed some day at startup differs from the normal information you see, it might mean that

- Someone has changed the system's normal BIOS configuration.
- The computer has reverted to default settings for troubleshooting or other reasons.
- The computer's battery is failing, causing stored setup information to be lost or corrupted.
- The hardware inside your computer has failed, or has been removed/replaced.

For more information on using the BIOS and CMOS setup to configure the system, correct configuration problems, or to determine the system's configuration, see Chapter 6 .

Hardware Resources

There are four types of hardware resources used by both onboard and add-on card devices:

- IRQ
- I/O port address
- DMA channel
- Memory address

Each device needs its own set of hardware resources, or needs to be a device that can share IRQs (the only one of the four resources that can be shared). Resource conflicts between devices can prevent your system from starting, lock up your system, or even cause data loss.

IRQs

IRQs, or interrupt requests, are a series of 8 or 16 lines that run between the CPU and both built-in and expansion card devices. Most devices use at least one IRQ. As the name implies, IRQs enable devices to interrupt the CPU to signal for attention when they need to send or receive data. Without IRQs, the CPU would run in isolation, never pausing to take care of any devices. The IRQ can be compared to a telephone's ringer. When the telephone rings, you pick it up. If the ringer is broken, you would never know if you were getting a telephone call. The different IRQs can be compared to a multiline phone in which each line is set aside for a different department. At least one IRQ is used by most major add-on cards (network, SCSI, sound, video, modem, and IEEE-1394) and major built-in system components such as ATA/IDE host adapters, serial, parallel, and USB ports.

tip

Many computers sold at retail stores do not display the POST and configuration information screens shown in Figures 2.24 and 2.25 as configured from the factory. To display this information, start the computer's BIOS setup program and look for a BIOS option called Quiet Boot. Disable this option and save your changes. Some systems have an option called Boot-Time Diagnostic Screen instead; enable this option and save your changes. When your computer restarts, it should display hardware information. Note that if your BIOS doesn't have an option such as Quiet Boot or Boot-Time Diagnostic Screen, you might not be able to view your hardware configuration at bootup.Because IRQs are used to handle different hardware devices, there's a simple rule of thumb: Generally, no two ISA devices should have the same IRQ assignment. While COM 1 and COM 3 serial ports are both assigned to IRQ 4, and COM 2 and COM 4 are both assigned to IRQ 3, this is not a true exception; if COM 1 and COM 3 are used at the same time, they stop working, and ditto for COM 2 and COM 4. PCI devices typically support IRQ sharing with Windows 95 OSR 2.x (Windows 95B), Windows 98, Windows Me, Windows 2000, and Windows XP, as well as Windows Server 2003 if the motherboard's design and configuration permit it.

The number of total IRQs available to you depends on the types of cards and expansion slots you use. Eight-bit ISA expansion slots, such as those used in the original IBM PC and PC/XT, could use IRQs up to 7 only. The ISA 16-bit and wider expansion slots can use up to IRQ 15. The typical and default IRQ usage appears in Table 2.2.

Table 2.2. Typical IRQ Usage

IRQStandard FunctionBus SlotResource TypeRecommended Use0System timerNoSystem—1Keyboard controllerNoSystem—2Second IRQ controller cascade to IRQ 9NoSystem—8Real-time clockNoSystem—9Available (might appear as IRQ 2)YesPCINetwork Interface Card or VGA[*]10AvailableYesPCIUSB[*]11AvailableYesPCISCSI host adapter[*]12Motherboard mouse port/availableYesISA/PCIMotherboard mouse port13Math coprocessorNoSystem—14Primary IDEYesPCIPrimary IDE (hard disks)15Secondary IDE/availableYesPCISecondary IDE (CD-ROM/tape)3Serial

Port 2 (COM 2:)YesISACOM 2:/internal modem4Serial Port 1 (COM 1:)YesISACOM 1:5Sound/Parallel Port 2 (LPT2:)YesISASound card[*]6Floppy disk controllerYesSystemFloppy controller7Parallel Port 1 (LPT1:)YesISALPT1:[*] On systems that use ACPI power management or support IRQ sharing, PCI versions of some or all of these devices might share IRQs.

IRQs are used by the system in the following way: You issue a command `DIR C:>LPT1` on the keyboard (IRQ 1). The IRQ is routed to the CPU, which "looks" at the keyboard and interprets the command. It uses IRQ 14 (IDE hard disk interface) to pass the command to the hard disk. Because the user has asked to print the command (`>LPT1`), the CPU uses IRQ 7 to "wake up" the parallel port and print the directory listing. Meantime, the system timer (IRQ 0) has kept everything working. As you look at the printout, you see the date- and timestamps next to the file listings, letting you know that the real-time clock (IRQ 8) is also on the job. Thus, even a simple operation uses several interrupt requests on several different lines.

Although Table 2.2 shows you traditional IRQ usage, you should realize that your computer might list much different IRQ usage and still work correctly. Here's why:

- If your computer is a so-called "legacy-free" system without serial, parallel, or PS/2 mouse and keyboard ports, or if you have manually disabled them, IRQs 3, 4, 7, and 12 will also be treated as PCI IRQs on most recent systems. This means they can be used for PCI cards not listed on the chart, such as IEEE-1394a host adapters, SCSI host adapters, video capture cards, add-on multi-I/O (serial/parallel) adapters, and so forth.
- Beginning with late versions of Windows 95 and continuing on to today's Windows XP versions, PCI devices can share IRQs on most systems. Although ISA devices such as built-in serial, parallel, and PS/2 mouse ports each need an exclusive IRQ, two or more PCI devices (as well as AGP video cards) can share IRQs as shown later in Figure 2.27.

Figure 2.27. The Windows XP Device Manager configured to display IRQ, DMA, and memory address usage.

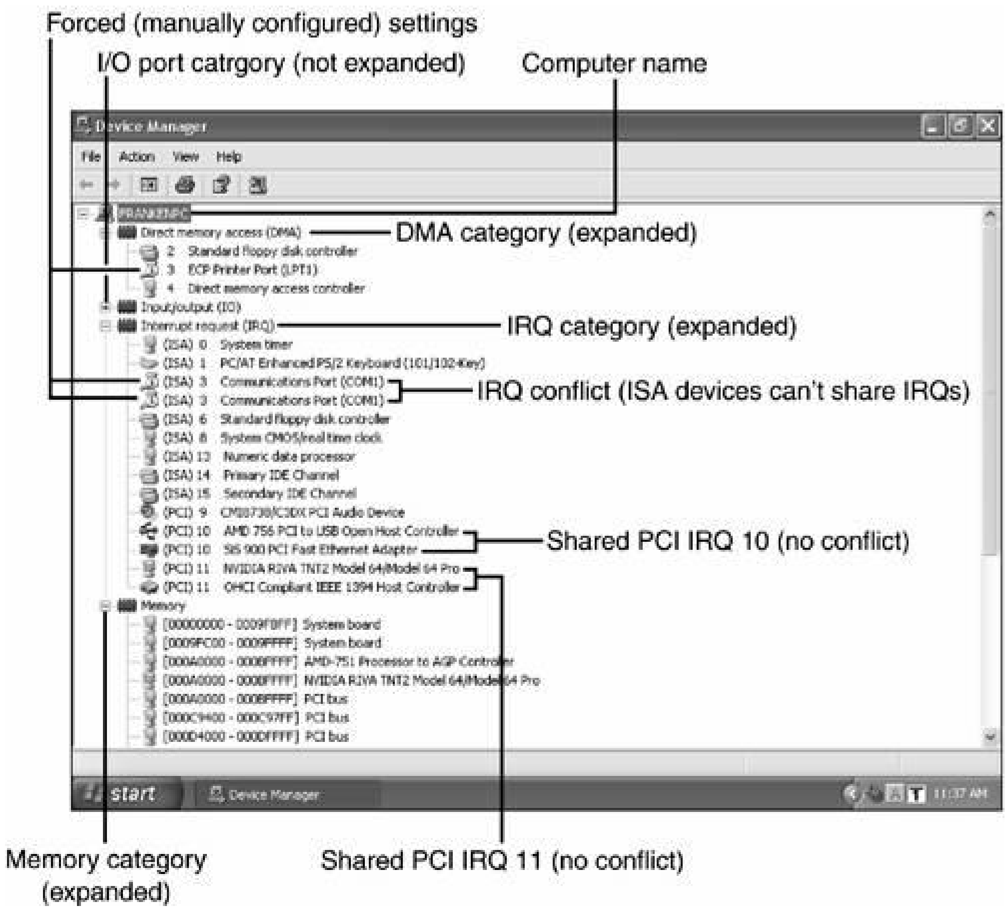

I/O Port Addresses

The system resource used even more often than IRQs, but one that causes fewer problems, is the **I/O port address** . The I/O port address is used to pass information between a given device and a system. Even simple devices that don't require an IRQ to function need one or more I/O port addresses. If you compare the IRQ to a telephone ringer, the I/O port address can be compared to the telephone transmitter and receiver, which do the actual work of sending and receiving your voice after you pick up the ringing telephone.

Your computer has 65,535 I/O port addresses (numbered hexadecimally), but uses only a relative few of them. Addresses might be followed by an "h", indicating hexadecimal. Each device uses an exclusive range of I/O port addresses for the actual transmission of data to and from the device. For situations in which two devices (such as COM 1 and COM 2) share a single IRQ, the I/O port addresses are still unique to each device. Often, information about an add-on card will list only the starting address, but you should know both the first and last I/O port address used by the device to keep from overlapping addresses, which will cause the devices to fail. There are thousands of I/O port addresses available in today's computers, so resource conflicts are rare unless a user tries to assign two serial ports to the same address (refer ahead to Figure 2.28).

Figure 2.28. The Windows XP Device Manager configured to display a portion of the I/O port address usage.

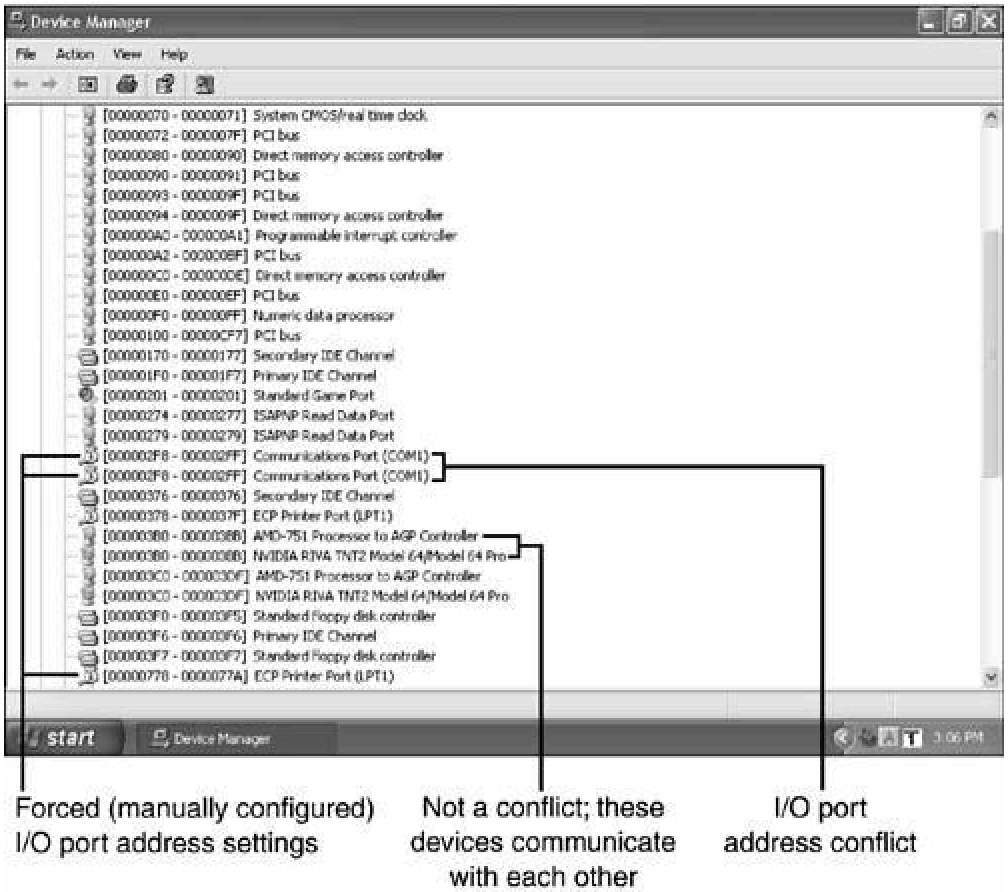

As you can see from Figure 2.26, different types of devices use different amounts of I/O address space. Some devices use a continuous "block" of addresses, such as the LPT1 (parallel) port, which uses eight addresses (0378-037Fhex). Others use only one or two addresses, or sometimes just a (literal) bit of a single address. For example, most systems don't have a secondary floppy controller, so the addresses reserved for it (0370–0375, bit 7 of 0377) will be available for another device.

Figure 2.26. A section of the I/O port address map appears here. Devices listed use the port addresses listed if they are present.

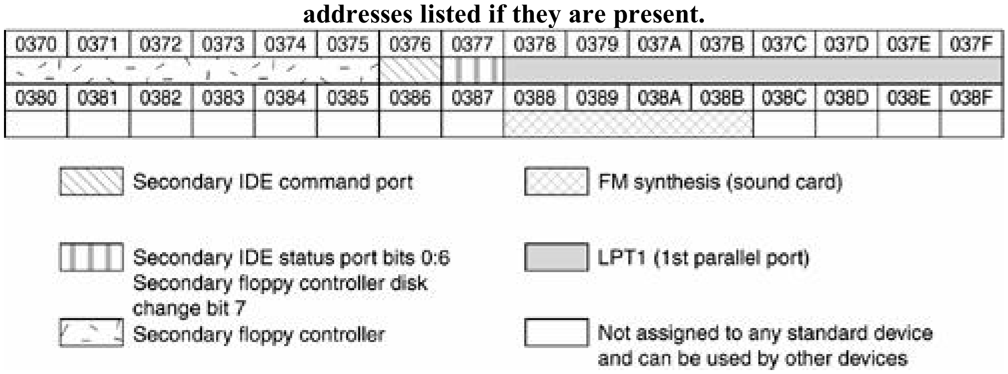

I/O port addresses are used by the system in the following way: You issue a command `DIR C:>LPT1` on the keyboard. The keystrokes pass through the keyboard interface's I/O addresses 0060h, 0064h. When the hard disk receives the command to read the directory listing, the command is given through I/O port 03F6, and the status is monitored through

bits 0:6 of I/O port 03F7. When the directory listing is redirected to the printer, the output is routed through LPT1's I/O port addresses 0378–037F.

As you look at I/O port address usage in your computer, you might see two different components, which are working but are displaying the same I/O port address per the Windows Device Manager. In these cases, the devices listed use the same I/O port address as a way to communicate with each other. For example, the AMD-751 chipset processor to AGP controller and the TNT2 AGP video card in my computer use the same I/O port addresses to facilitate communication with each other as shown in Figure 2.27.

DMA (Direct Memory Access) Channels

As you saw earlier in this chapter, the CPU is responsible for many tasks, including that of being a sort of "traffic cop," overseeing the transfer of information between itself, memory, and various devices. Although a police officer on the corner of a busy intersection helps keep bumper-to-bumper rush-hour drivers on their best behavior, this manual stop-go-stop process isn't the fastest way to get around town—or around the motherboard.

The expressway, beltway, or freeway that bypasses surface streets, traffic lights, and hand-signaling police officers is a faster way to travel when conditions are favorable. Similarly, bypassing the CPU for memory to add-on board transfers (either direction) is also a faster way to travel.

This process of bypassing the CPU is called Direct Memory Access (DMA). DMA transfers can be done in two ways; some DMA transfers (such as those done by tape backup drives or by PCI cards using bus-mastering) do not require a particular DMA channel; however, some devices, such as popular ISA sound cards and the ECP mode of the parallel port, require that we select an unused DMA channel. Table 2.3 lists the standard DMA channel uses. In Table 2.3, standard uses appear in bold type, and typical uses (which can change) are shown in italic type.

Table 2.3. Standard and Typical DMA Channel Uses

DMA Channel #	Use	Notes
0	Some sound cards	Requires 16-bit ISA or PCI card[*]
1	Sound card	Sound Blaster standard; also used by "clones"
2	Floppy drive	
3	LPT port in ECP mode	Some systems offer DMA 1 as an alternative setting
4	System reserved	
5	Sound card	Requires 16-bit ISA or PCI card[*]; some sound cards use only DMA 1 or use DMA 0 instead of 5
6		Requires 16-bit ISA card
7		Requires 16-bit ISA card

[*] PCI sound cards use DMA only if used in a Sound Blaster emulation mode.

Because DMA is used for high-speed data transfer, and because there are relatively few DMA channels, some users are tempted to "share" them between ISA devices. Never do this. If two devices using DMA are used at the same time, a catastrophic loss of data could result. Because the CPU is not involved in DMA transfers, there's no "traffic cop" in case of disaster!

One of my favorite "disaster" stories was related by longtime Byte magazine columnist Jerry Pournelle in July 1994. He installed a sound card on a system that already had an unusual hard disk interface that used DMA; both devices were set to DMA 5. When he

tried to record sound samples with the sound card, the flow of hard disk data and the flow of sampled sound "collided" on DMA 5, wiping out the entire contents of his hard disk! DMA conflicts are rare, especially now that ISA cards are not used by most systems, but they're never funny.

Memory Addresses

The physical memory (RAM) installed into the computer is divided into memory addresses; each address equals a byte of RAM. For normal operations, the system automatically determines which memory address to use for retrieving existing information or for temporary location of new memory addresses. These addresses are also given in hexadecimal notation.

Like DMA channels, memory addresses are also abundant in computers. Add-on cards that have their own BIOS chips (some SCSI, some ATA/IDE, and some network cards, as well as all VGA and 3D video cards) or add-on cards that have RAM (video cards and a few network cards), however, must use unique memory addresses that are found in the range between 640 kilobytes (KB) and 1 megabyte (MB). Because there's abundant memory address space and relatively low demand for memory addresses, conflicts are rare unless you manually configured a card to use an address already in use by another device. Memory managers such as EMM386.EXE that can be used to support MS-DOS software running under Windows 9x can also use memory addresses that are not in use by cards.

Viewing Hardware Resources in Use

You can view the current resource usage in your computer with the Windows Device Manager. To see the resource usage for a particular device, open the Device Manager, open the device's properties sheet, and click Resources. You can also use the Windows System Information program to view resource usage. For details, see Chapter 18 , "Using and Optimizing Windows."

If you need to install non-PnP devices, or if you are concerned about installing devices that have limited configuration options, you can also view all the resources currently in use. With Windows XP and Windows 2000, start the Device Manager, select View, Resources by Type, and click the plus sign (+) next to each category. Figure 2.27 shows the DMA channel, IRQ, and a portion of the memory resource usage in one of my computers, which runs Windows XP. Figure 2.28 shows a portion of the I/O port address usage in the same computer. With Windows 9x/Me, double-click the computer icon at the top of the Device Manager listing to see resource usage.

Note that in these two figures, forced (manually configured) settings are indicated with a white circle containing a blue . Forced settings are seldom a good idea (plug-and-play configuration usually works much better, especially with Windows XP). Additionally, if two ISA devices have been forced to use the same hardware resource (such as the COM ports in Figures 2.27 and 2.28, which are both set to use the same IRQ and I/O port address), you have a hardware conflict that will prevent the devices from working.

Study Lab

Don't miss the Study Lab materials found on the CD accompanying this book. Each Study Lab is tailored to the individual chapters in this book, meaning that you'll quickly be able to determine which topics you understand well enough to pass the exam and which topics

need more study. The Study Labs are presented in printable PDF format so that you can take them with you to study at work, on the road, or even in your car just before test time!

The Absolute Minimum

Here are the high points of this chapter. Read these just before you take your exam, or any time to find out what topics you need to review in more detail. All computers contain hardware, software, and firmware. There are two types of memory RAM (contents change as programs and data are loaded, created, and saved) and ROM (contents can't be changed under normal circumstances). Hardware subsystems such as video, drives, memory, and others are controlled by the system BIOS on the motherboard and by the operating system and its device drivers. The outside of the computer typically has cable connections for peripherals, drive bays for removable-media drives, the power supply, and the power switch. The inside of the computer contains the motherboard, add-on cards, memory modules, the processor, and internal drives (hard disk, floppy disk, and optical drives). The BIOS controls built-in components and displays the installed components on some systems. IRQs range from 0–15 and are used to request attention from the CPU. I/O port addresses are used to transmit information between the system and a particular device. There are 65,535 I/O port addresses. DMA channels are used by ISA devices to bypass the processor for faster device-memory transfers. DMA channels range from 0 to 7. Memory addresses are used by devices with built-in ROM or RAM chips. The Windows Device Manager and Windows System Information display current hardware (IRQ, DMA, I/O port address, and memory address) usage.